arXiv:1604.00111v2 [cs.DB] 7 Apr 2016

Computing Join Queries with Functional Dependencies

Mahmoud Abo Khamis
LogicBlox, Inc.
& SUNY Buffalo

Hung Q. Ngo
LogicBlox, Inc.
& SUNY Buffalo

Dan Suciu
LogicBlox, Inc.
& University of Washington

Abstract

Recently, Gottlob, Lee, Valiant, and Valiant (GLVV) presented an output size bound for join queries with functional dependencies (FD), based on a linear program on polymatroids. GLVV bound strictly generalizes the bound of Atserias, Grohe and Marx (AGM) for queries with no FD, in which case there are known algorithms running within AGM bound and thus are worst-case optimal.

A main result of this paper is an algorithm for computing join queries with FDs, running within GLVV bound up to a poly-log factor. In particular, our algorithm is worst-case optimal for any query where the GLVV bound is tight. As an unexpected by-product, our algorithm manages to solve a harder problem, where (some) input relations may have prescribed maximum degree bounds, of which *both* functional dependencies and cardinality bounds are special cases.

We extend Gottlob et al. framework by replacing all variable subsets with the lattice of closed sets (under the given FDs). This gives us new insights into the structure of the worst-case bound and worst-case instances. While it is still open whether GLVV bound is tight in general, we show that it is tight on distributive lattices and some other simple lattices. Distributive lattices capture a strict superset of queries with no FD and with simple FDs. We also present two simpler algorithms which are also worst-case optimal on distributive lattices within a single-log factor, but they do not match GLVV bound on a general lattice. Our algorithms are designed based on a novel principle: we turn a proof of a polymatroid-based output size bound into an algorithm.

1 Introduction

Several results published in the last ten years or so have lead to tight worst-case output size bounds (*AGM-bound* [3, 13]) and the development of a new class of query processing algorithms running within the bound's time budget [18,19,23]. These new *worst-case optimal* algorithms are quite different from traditional query plans, in the sense that they no longer compute one pairwise join at a time, but instead process the query globally. The runtime is bounded by $O(N^{\rho^*})$, where N is the size of the database, while ρ^* is the value of the optimal fractional edge cover of the query. For example, they compute the triangle query $Q(x, y, z)$:- $R(x, y), S(y, z), T(z, x)$, in worst-case time $O(N^{3/2})$, while any traditional query plan requires worst-case time $\Omega(N^2)$ [19].

While the vast majority of database engines today still rely on traditional query plans, new, complex data analytics engines increasingly switch to worst-case optimal algorithms: LogicBlox' engine [2] is built on a worst-case optimal algorithm called LeapFrog Triejoin [23] (LFTJ), the *Myria* data analytics platform supports a variant of LFTJ [7], and the increased role of SIMD instructions in modern processors also favors the new class of worst-case optimal algorithms [1].

When there are functional dependencies (FDs), however, the AGM-bound is no-longer tight. Gottlob, Lee, Valiant, and Valiant [11] initiated the study of the worst-case output size of a query in the presence of FDs and described an upper bound (reviewed in Section 2), called the *GLVV-bound* henceforth. It remains open whether the GLVV bound is tight; at the present, it is the best known upper bound for the query size in the presence of FDs. A recent result by Gogacz and Toruńczyk [10] proved a weaker statement, namely that if we replace the polymatroidal constraints in the GLVV-bound with entropic constraints (of which there are infinitely many), then the bound is tight.

Our paper proposes a novel approach to studying queries in the presence of FDs, by using lattice theory. We present several theoretical results that clarify precisely where queries with FDs become more difficult to study than those without. We describe a an algorithm which runs in time proportional to the GLVV-bound, within a polylogarithmic factor. We also describe two special cases where the polylogarithmic factor is reduced to a single log. Before discussing our results, we explain their significance.

1.1 Motivation

User-Defined Functions. UDF's, or interpreted predicates, significantly affect the runtime of a query. Consider:

$$Q(x, y, z, u) \text{ :- } R(x, y), S(y, z), T(z, u), u = f(x, z), x = g(y, u). \tag{1}$$

The predicates $u = f(x, z)$ and $x = g(y, u)$ represent two user-defined functions, f and g; for example $f(x, z)$ could be $x + z$, or the concatenation of two strings x and z. Any UDF can be modeled as a relation with a primary key, for example the function f can be viewed as a relation $F(x, z, u)$ of cardinality N^2 (with one entry for every x, z pair) satisfying the FD $xz \to u$; similarly g introduces the FD $yu \to x$. The addition of these two FD's significantly affects the output size and the query evaluation complexity: if we first computed the intermediate query $R(x, y), S(y, z), T(z, u)$ then applied the two predicates $u = f(x, z)$ and $x = g(y, u)$ then the runtime can be as high as N^2 because the size of the intermediate query is N^2 when $R = \{(i, 1) \mid i \in [N]\}$, $S = \{(1, 1)\}$, $T = \{(1, i) \mid i \in [N]\}$. However, as we will show, the GLVV bound for the size of Q is $\leqslant N^{3/2}$. Our new algorithms run in $\tilde{O}(N^{3/2})$, and thus reduce the running time asymptotically. As shall be seen, this runtime is also worst-case optimal.

A related problem is querying relations with restricted access patterns [5]. In that setting, some of the relations in the database can only be read by providing the values of one or more attributes. As shown above, a user defined function $f(x, z)$ can be modeled as an infinite relation $F(u, x, z) \equiv (u = f(x, z))$ with the restriction that F can be accessed only by providing inputs for the variables x and z. The work on querying under restricted access patterns has focused on the answerability question (whether a query can or cannot be answered). Our work extends that, by finding algorithms for answering the query within GLVV-bound.

Known Frequencies. Systems often know an upper bound on the frequencies (or degrees) in the database. For example, consider the triangle query above, and assume that the binary graph defined by the relation $R(x, y)$ has a bounded degree: all outdegrees are $\deg(x) \leqslant d_1$ and all indegrees are $\deg(y) \leqslant d_2$. One can model this scenario by introducing an artificial color c_1 on outgoing edges, and a color c_2 for incoming edges of R:

$$\begin{aligned} Q(x, y, z, c_1, c_2) \text{ :- } & R(x, c_1, c_2, y), S(y, z), T(z, x), C_1(c_1), C_2(c_2), \\ & xc_1 \to y, yc_2 \to x, xy \to c_1c_2 \end{aligned} \tag{2}$$

where $|R| = |S| = |T| = N$ and $|C_1| = d_1, |C_2| = d_2$. Thus, each outgoing edge from x is colored with some distinct color c_1, and similarly each incoming edge to y is colored with some distinct color c_2. The new predicates C_1, C_2 limit the number of colors to d_1, d_2 respectively. We will show that the worst case query output decreases from $N^{3/2}$ to $\min(N^{3/2}, Nd_1, Nd_2)$ Alternatively, we can use the linear program and algorithm in Sec 5.3 to capture queries with known maximum degree bounds. We note that another approach to evaluate a limited class of queries over databases with known degrees has been recently described in [15].

1.2 Overview of the Results

Grohe and Marx [13] and later Atserias, Grohe and Marx [3] derived an elegant tight upper bound on the output size of a join query: $\prod_{j=1}^{m} |R_j|^{w_j}$, where $R_1, \ldots, R_m$ are the input relations to the query, and $(w_j)_{j=1}^{m}$ is any fractional edge cover of the query's hypergraph. This bound is known today as the AGM bound. A simple example that gives a great intuition for this formula, due to Grohe [12], is the following. Suppose we choose a subset of relations $R_{j_1}, R_{j_2}, \ldots$ that together contain all variables of the query, in other words they form an integral cover of the query's hypergraph. Then, obviously $|Q|$ is upper bounded by the product $|R_{j_1}||R_{j_2}| \cdots$, because the output to Q is contained in the cross product of these relations. The AGM bound simply generalizes this property from integral to fractional edge covers. They proved that the bound is tight by describing a simple database instance for any query, such that the query's output matches the bound. In that instance, every relation is a cross product of sets, one set per variable; we call it a *product instance*.

An obvious open question was whether a query Q can be evaluated on any database instance D in time that is no larger than the AGM bound of Q on databases with the same cardinalities as D; such an algorithm is called *worst-case optimal*. Ngo, Porat, Ré, and Rudra [18] described the first worst-case optimal algorithm; later Veldhuizen [23] proved that LFTJ, an algorithm already implemented at LogicBlox earlier, is also worst-case optimal. A survey and unification of these two algorithms can be found in [19].

Neither the AGM bound nor the associated algorithms analytically exploit FDs in the database.[1] Such FDs can provably reduce the worst-case output of a query, but the upper bound and algorithms mentioned above cannot use this information, and instead treat the query by ignoring the FDs. Gottlob et al. [11] studied the upper bound of the query size in the presence of FDs, and established two classes of results. The first was a characterization of this bound in the case when the FD's are restricted to simple keys; as we will show, this case can be solved entirely using the AGM bound by simply replacing each relation with its closure. Next, they described a novel approach to reasoning about the worst-case output of a query, using information theory. They viewed the query output as a multivariate probability space, and introduced two constraints on the marginal entropies: a cardinality constraint for each input relation R, stating that the entropy of its variables cannot exceed the uniform entropy $H(\text{vars}(R)) \leqslant \log_2 |R|$, and one constraint for each FD $X \to Y$, stating $H(XY) = H(X)$. (Note that X and Y are sets of variables.) The largest answer to the query Q is then given by the largest possible value of $2^{H(\text{vars}(Q))}$, over all choices of entropic functions H that satisfy these constraints. But characterizing the space of all entropic functions H is a long standing open problem in information theory [24]; to circumvent that, they relax the function H by allowing it to be any function that satisfies Shannon inequalities. Such a function is called a *polymatroid* in the literature, and we denote it with lower case h to distinguish it from entropic functions H. Thus, the problem in [11] can be

[1] Algorithmically, LFTJ handles FDs by binding variables at the earliest trie level at which they are functionally determined. For example, in $R(x, y), S(y, z), T(z, u), u = f(x, z), x = g(y, u)$ with a key order $[x, y, z, u]$, whenever z was bound the value of u would be immediately computed by $u = f(x, z)$.

stated equivalently as: find the maximum value $h(\text{vars}(Q))$ where h ranges over all polymatroids satisfying the given constraints.

In this paper, we continue the study of query evaluation under general FDs. We establish both bounds and algorithms. Our novelty is to model FDs as a lattice $\mathbf{L}$ of the closed sets of attributes, and to study polymatroids on lattices. The function h is now any non-negative, monotone, sub-modular function (i.e. $h(X) + h(Y) \geqslant h(X \vee Y) + h(X \wedge Y)$) that satisfies all cardinality constraints. FD constraints are enforced *automatically* by the lattice structure and the upper bound on the query size is $2^{\max h(\text{vars}(Q))}$. When there are no FDs, the lattice is a Boolean algebra.

Our first question is whether the elegant AGM bound and worst-case product instance carries over to arbitrary FDs. We answer this question completely, by proving that both sides of the AGM bound hold iff the lattice has a special structure, which we call a *normal lattice*. Both upper and lower bounds require minor extensions to be applicable to normal lattices. The standard AGM upper bound is given in terms of fractional edge covers of the query's hypergraph, but in a normal lattice one needs to consider a dual hypergraph, whose nodes are L's co-atoms. In a Boolean algebra, these two hypergraphs are isomorphic, because of the bijection $X \mapsto (\text{vars}(Q) - \{X\})$ between variables and co-atoms, but in a general lattice they can be significantly different. The notion of normal lattice seems novel, and strictly includes all distributive lattices, which in turn include all lattices corresponding to simple FDs (each FD is of the form $a \to b$, where a, b are attributes). The second minor change is that one needs to allow for a slight generalization of product instances, to what we call quasi-product instances. Importantly, both these properties fail on non-normal lattices; in particular, worst-case instances cannot be quasi-free.

The canonical example of a non-normal lattice is M_3 (one of the two canonical non-distributive lattices, see the right part of Fig. 3). Every lattice L that contains M_3 as a sublattice such that $\max L = \max M_3$ is non-normal; we conjecture that the converse also holds. Interestingly, the other canonical non-distributive lattice N_5 is normal.

Next, we examine algorithms whose runtime is bounded by the GLVV bound. We propose a novel methodology for designing such algorithms, starting from the observation that such an algorithm must provide a proof of the query's upper bound, equivalently, a proof of an inequality of the form $\sum_j w_j h(\text{vars}(R_j)) \geqslant h(\text{vars}(Q))$, where $R_1, R_2, \ldots$ are the input relations. In the case of a Boolean algebra, Shearer's lemma [8] is of this form; in a normal lattice this corresponds to a fractional edge cover of the co-atomic hypergraph; and, for a general lattice it is a general inequality. A key motivation behind NPRR [18] was to prove inequalities algorithmically. This paper completes the cycle by proceeding in the opposite direction: given a proof method for such inequalities, we design algorithms whose steps correspond to the proof steps. We design three such algorithms, corresponding to three methods for proving the above type of inequalities.

The first algorithm called the *chain algorithm* runs within the *chain bound*; the proof technique is adapted from Radhakrishnan's proof [20] of Shearer's lemma to general lattices. Both the chain bound and algorithm strictly generalize AGM-bound and worst-case optimal algorithms for join queries *without* FDs. The second algorithm, called the *sub-modular algorithm*, runs within the *sub-modularity bound*; the proof technique is that of Balister and Bollobás's [4]. The third algorithm, called the *conditional sub-modularity algorithm* (CSMA) runs within the general *GLVV bound*, up to a polylogarithmic factor; the proof technique is our own, based on linear programming duality. In addition to being able to achieve the most general bound, CSMA can be used straightforwardly to handle input relations with known maximum degree bounds. We remark that GLVV bound is stronger than both the other two bounds. We show that they are tight, and thus our algorithms are worst-case optimal, for distributive lattices.

Outline. The paper is organized as follows. Background material is reviewed in Sec. 2, and basic definitions for our lattice-based approach are given in Sec. 3. We describe the main result on normal lattices in Sec. 4, then present our three algorithms and bounds in Sec. 5.

2 Notations and Prior Results

For any positive integer n, $[n]$ denotes the set $\{1, \ldots, n\}$. All log in the paper are of base 2. We fix a relational schema $\mathbf{R} = \{R_1, \ldots, R_m\}$ whose attributes belong to a set of attributes $\mathbf{X} = \{x_1, \ldots, x_k\}$. We refer to $x_i \in \mathbf{X}$ interchangeably as an *attribute* or a *variable*; similarly we refer to R_j as a *relation*, or an *input*. We use lower case letters $x \in \mathbf{X}$ to denote single variables, and upper case letters $X \subseteq \mathbf{X}$ to denote sets of variables. The *domain* of variable x is denoted by $\mathsf{Domain}(x)$. For each relation R_j we denote $\text{vars}(R_j) = X_j \subseteq \mathbf{X}$ its set of attributes, and sometimes blur the distinction between R_j and X_j, writing, with some abuse, $R_j \subseteq \mathbf{X}$. We consider full conjunctive queries without self-joins:

$$Q(x_1, \ldots, x_k) \coloneq R_1(X_1), \ldots, R_m(X_m) \tag{3}$$

We will drop variables from the head, since it is understood that all variables need to be listed.

A database instance D consists of one relational instance R_j^D for each relation symbol; we denote $N_j = |R_j^D|$, $N = |D| = \sum_j N_j$, and use lower case for logs, $n_j = \log_2 N_j$. We denote Q^D the answer to Q on the database instance D. A *product database instance* is a database instance such that $R_j^D = \prod_{x_i \in R_j} \mathsf{Domain}(x_i)$ for $j \in [m]$; the query answer on a product database is the cross product of all domains, $Q^D = \prod_{i=1}^k \mathsf{Domain}(x_i)$.

The AGM Bound. A series of results over the last ten years [3, 13, 18, 19, 23] have established tight connections between the maximum output size of a query and the runtime of a query evaluation algorithm. The *query hypergraph* of a query Q is $H_Q = (\mathbf{X}, \mathbf{R})$; its nodes are the variables and its hyperedges are the input relations (where each $R_j \in \mathbf{R}$ is viewed as a set of variables). Consider the following two linear programs (LP's), called weighted fractional edge cover LP and vertex packing LP, respectively:

(Weighted) Fractional Edge Cover

$$\begin{array}{lrcll} \text{minimize} & \sum_j w_j n_j & & & \\ & \sum_{j: x_i \in R_j} w_j & \geqslant & 1, & \forall i \in [k] \\ & w_j & \geqslant & 0, & \forall j \in [m] \end{array}$$

(Weighted) Fractional Vertex Packing

$$\begin{array}{lrcll} \text{maximize} & \sum_i v_i & & & \\ & \sum_{i: x_i \in R_j} v_i & \leqslant & n_j, & \forall j \in [m] \\ & v_i & \geqslant & 0, & \forall i \in [k] \end{array}$$

We call a feasible solution to the first LP a *fractional edge cover*, and to the second a *fractional vertex packing*. The traditional (unweighted) notions correspond to $n_j = 1, \forall j$.

Theorem 2.1 (AGM bound)**.** *[3, 13] (1) Let $(w_j)_{j=1}^m$ be a fractional edge cover. Then, for any input database D s.t. $|R_j^D| \leqslant N_j$ for all $j \in [m]$, the output size of Q is bounded by $2^{\sum_j w_j n_j}$. In other words, $|Q^D| \leqslant \prod_j N_j^{w_j}$. (2) Let $v_i, i \in [k]$ be a fractional vertex packing, and let D be the product database instance where $D_i = [2^{v_i}]$. Then, $|Q^D| = \prod_i 2^{v_i}$.*

(In the statement above, we ignore the issue of integrality of the v_i for the sake of clarity. There is a bit of loss when v_i are not integers, but this fact does not affect the asymptotics of the lowerbound [3].) By strong duality, these two LPs have the same optimal objective value, denoted by $\rho^*(Q, (N_j)_{j=1}^m)$ [12]. Let $(w_j^*)_{j=1}^m$ be the optimal edge cover. The *AGM bound* of the query Q is

$$\mathsf{AGM}(Q, (N_j)_{j=1}^m) = 2^{\rho^*(Q, (N_j)_{j=1}^m)} = \prod_j N_j^{w_j^*},$$

or just $\mathsf{AGM}(Q)$ when the cardinalities $(N_j)_{j=1}^m$ are clear from the context. The query's output size is always $\leqslant \mathsf{AGM}(Q)$, and this bound is tight, because on the product database described above, the output is $\mathsf{AGM}(Q)$. It is easy to check that $\mathsf{AGM}(Q) = \min_w \prod_j N_j^{w_j}$, where w ranges over the vertices of the edge cover polytope. For example, for $Q = R(x, y), S(y, z), T(z, x)$ the edge cover polytope has vertices $\left\{(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}), (1, 1, 0), (1, 0, 1), (0, 1, 1)\right\}$, thus,

$$\mathsf{AGM}(Q) = \min(\sqrt{N_R N_S N_T}, N_R N_S, N_R N_T, N_S N_T) \tag{4}$$

Several query evaluation algorithms have been described in the literature with runtime[2] $\tilde{O}(N + \mathsf{AGM}(Q))$: NPRR [18], LFTJ [23], Generic-join [19].

Functional Dependencies. A query with functional dependencies is a pair (Q, FD), where Q is a query and FD is a set of *functional dependencies* (fd), which are expressions of the form $U \to V$ where $U, V \subseteq \mathbf{X}$. An fd can be either defined by some relation $R_j(X_j)$, in which case we call it *guarded* (in particular $U, V \subseteq X_j$), or can be defined by a UDF (as we saw in Sec. 1.1), and then we call it *unguarded*. A *simple* fd is of the form $u \to v$ where both u, v are variables, and a *simple key* is a simple fd guarded in some R_j s.t. u is a key for R_j.

Output size bounds in the presence of a set of fd's was studied in [11]. Their bound is defined only in terms of the maximum cardinality, $\max_j N_j$, but in this paper we generalize the discussion to all cardinalities $(N_j)_j$. The key technique introduced in [11] consists of using information theory to upper bound the size of the query, as reviewed next.

Let Q^D be the query answer over some instance D. Define a probability distribution over $\prod_{i=1}^k D_i$ by randomly drawing one tuple from Q^D with probability $1/|Q^D|$ each. Under this distribution, the joint entropy of the k random variables $\mathbf{X}$ is $H(\mathbf{X}) = \log_2 |Q^D|$. Each subset of variables $X \subseteq \mathbf{X}$ defines a (marginal) distribution, with entropy $H(X)$. If $X = X_j$ $(= \text{vars}(R_j))$, then H must satisfy the following *cardinality constraint* $H(X_j) \leqslant \log_2 N_j = n_j$, because $\Pi_{X_j}(Q^D) \subseteq R_j^D$ and the marginal entropy is bounded above by the uniform marginal entropy. With some abuse we write $H(R_j) \leqslant \log_2 N_j$, blurring the distinction between R_j and X_j. In addition, for any fd $U \to V$ the entropy must satisfy the *fd-constraint* $H(U) = H(UV)$.

We give here a very simple illustration of how the approach in [11] models the query output using entropy, by illustrating on the query $Q(x, y, z) = R(x, y), S(y, z), T(z, x)$ (without fd's) and output with five outcomes:

x	y	z	
a	3	r	1/5
a	2	q	1/5
b	2	q	1/5
d	3	r	1/5
a	3	q	1/5

x	y	
a	3	2/5
a	2	1/5
b	2	1/5
d	3	1/5

y	z	
3	r	2/5
2	q	2/5
3	q	1/5
4	q	0

x	z	
a	r	1/5
a	q	2/5
b	q	1/5
d	r	1/5

Here $H(xyz) = \log 5$, $H(xy) \leqslant \log |R| = \log 4$, $H(yz) \leqslant \log |S| = \log 4$, and $H(xz) \leqslant \log |T| = \log 4$.

GLVV Bound. Gottlob et al. [11] observe that $\log |Q^D| \leqslant \max_H H(\mathbf{X})$, where H ranges over all entropic functions that satisfy the cardinality constraints and the fd-constraints, and this bound has recently been shown to be tight [10]. However, computing this upper bound is extremely difficult, because of a long standing open problem in information theory: the characterization of the cone of the closure of the set of entropic functions. To circumvent this difficulty, in [11] the entropic function H is replaced with a polymatroid function h. A *polymatroid* over a set of variables $\mathbf{X}$ is a function $h : 2^{\mathbf{X}} \to \mathbb{R}^+$ that satisfies the following inequalities, called *Shannon inequalities*:

$$\begin{aligned} h(X) + h(Y) &\geqslant h(X \cap Y) + h(X \cup Y) && \text{Sub-modularity} \\ h(X \cup Y) &\geqslant h(X) && \text{Monotonicity} \\ h(\varnothing) &= 0 && \text{Zero} \end{aligned}$$

Every entropic function H is a polymatroid, and the converse fails when $\mathbf{X}$ has four or more variables [25]. The bound introduced in [11] is $\mathsf{GLVV}(Q, \mathsf{FD}, (N_j)_j) \stackrel{\text{def}}{=} \max_h h(\mathbf{X})$, where h ranges over all polymatroids that satisfy all cardinality constraints $h(X_j) \leqslant n_j$, and all fd-constraints $h(UV) = h(U)$ for $U \to V \in \mathsf{FD}$. We abbreviate it $\mathsf{GLVV}(Q)$ when FD and $(N_j)_j$ are clear from the context. Clearly $\mathsf{GLVV}(Q)$ is an upper bound on $|Q^D|$, and it is open whether the bound is always tight.

[2] $\tilde{O}$ means up to a logarithmic factor.

Closure. Fix a set FD. The *closure* of a set $X \subseteq \mathbf{X}$ is the smallest set X^+ such that $X \subseteq X^+$, and that $U \to V \in \mathsf{FD}$ and $U \subseteq X^+$ imply $V \subseteq X^+$. The intersection $X^+ \cap Y^+$ of closed sets is closed, hence the family $\{X^+ \mid X \subseteq \mathbf{X}\}$ is a *closure*.

Denote by Q^+ the query obtained by replacing each relation $R_j(X_j)$ with $R_j(X_j^+)$, then forgetting all functional dependencies; it is easy to check that $\max_D |Q^D| \leqslant \mathsf{AGM}(Q^+)$. This bound is tight when all fd's are simple keys, because any product database over the schema of Q^+ can be converted into a database over the schema for Q that satisfies all simple keys. Thus, for simple keys, $\mathsf{AGM}(Q^+)$ is a tight upper bound on $|Q^D|$. Theorem 4.4 in [11] uses query coloring to prove essentially the same result. For a simple illustration, consider $Q :- R(x,y), S(y,z), T(z,u), K(u,x)$, where $\mathsf{AGM}(Q) = \min(|R| \cdot |T|, |S| \cdot |K|)$. If we define y to be a key in S, in other words $\mathsf{FD} = \{y \to z\}$, then $Q^+ = R(x,y,z), S(y,z), T(z,u), K(u,x)$ and $AGM(Q^+) = \min\{|R| \cdot |T|, |S| \cdot |K|, |R| \cdot |K|\}$. However, for fd's other than simple keys this technique fails, as illustrated by $Q(x,y,z) = R(x), S(y), T(x,y,z)$ where xy is a key in T (in other words $xy \to z$): when $|R| = |S| = N$, $|T| = M \gg N^2$ then $Q^+ = Q$ and $\mathsf{AGM}(Q^+) = M$, yet one can easily verify that $|Q^D| \leqslant N^2$.

The Expansion Procedure. All our algorithms use the following simple subroutine, called an expansion. Fix a relation $R(X)$ and a database instance D. R may be an input relation, or some intermediate relation generated during query evaluation. An *expansion* of R^D is some relation $(R^+)^D$ over attributes X^+ such that $\Pi_{X^+}(Q^D) \subseteq (R^+)^D$ and $\Pi_X((R^+)^D) \subseteq R^D$. If $X^+ = X$, then the expansion could be R^D itself, or any partial semi-join reduction that removes dangling tuples from R^D (which do not join with tuples in other relations). If $X \neq X^+$, then the expansion fills in the extra attributes, by repeatedly applying functional dependencies $X \to y$: if the fd is guarded in R_j, then it joins R with $\Pi_{Xy}(R_j)$; otherwise, if the fd corresponds to a UDF $y = f(X)$ then it simply computes f for each tuple in R. The expansion of R can be done in time $\tilde{O}(N)$ using standard techniques.

3 A Lattice-based Approach

3.1 Lattice representation of queries with FDs

Fix a query with functional dependencies (Q, FD), over variables $\mathbf{X}$. It is well known [9, 14, 16] that the set of closed sets forms a lattice:

Definition 3.1 (Lattice representation)**.** The lattice associated to FD is $\mathbf{L}_{\mathsf{FD}} = (L, \preceq)$, where L consists of all closed sets and the partial order $\preceq$ is $\subseteq$. The *lattice representation* of a query (Q, FD) is the pair $(\mathbf{L}_{\mathsf{FD}}, \mathbf{R}^+)$, where $\mathbf{R}^+ = \{R_1^+, \ldots, R_m^+\} \subseteq L$ is the set of closures of the attributes of the input relations. We drop the subscript FD when it is clear from the context and write simply $\mathbf{L}$. With the expansion procedure in place, w.l.o.g. we assume that all the input R_j are closed sets; thus, the lattice representation of the query can be denoted simply by $(\mathbf{L}, \mathbf{R})$, where $\mathbf{R} = \{R_1, \ldots, R_m\} \subseteq L$, and $\bigvee \mathbf{R} = \bigvee_{R_j \in \mathbf{R}} R_j = \hat{1}$.

Note that the size of the lattice may be exponential in that of the query, but this does not affect our complexity results, because they are given in terms of data complexity only. If Q has no functional dependencies, then $\mathbf{L}$ is the Boolean algebra $2^{\mathbf{X}}$. We will use Fig.1 as a running example, which illustrates the lattice for query (1): the lattice elements corresponding to input relations are framed.

We briefly review some notions of lattice theory needed later and refer the reader to [22] for an extensive treatment. Let $\wedge, \vee, \hat{0}, \hat{1}$ denote the greatest lower bound (*meet*), least upper bound (*join*), minimum and maximum elements of the lattice. The following hold in $\mathbf{L}_{\mathsf{FD}}$: $X \wedge Y = X \cap Y$, $X \vee Y = (X \cup Y)^+$, $\hat{0} = \varnothing$, and $\hat{1} = \mathbf{X}$. We write $X \lessgtr Y$ to mean X and Y are incomparable. An element U is said to *cover* an element V of L if $U \succ V$ and $U \succeq W \succeq V$ implies $W = U$ or $W = V$; an *atom* is an element X that covers $\hat{0}$; a *co-atom* is an element covered by $\hat{1}$; X is called a *join-irreducible* if $Y \vee Z = X$ implies $Y = X$ or $Z = X$; X is called *meet-irreducible* if $Y \wedge Z = X$ implies $Y = X$ or $Z = X$. For any $X \in L$, let $\Lambda_X \stackrel{\text{def}}{=} \{Z \mid Z \preceq X, Z \text{ is a join-irreducible}\}$. The mapping $X \mapsto \Lambda_X$ defines a 1-1 mapping from L to the order ideals of the poset of join-irreducibles, with inverse $\Lambda_X \mapsto \bigvee \Lambda_X = X$.

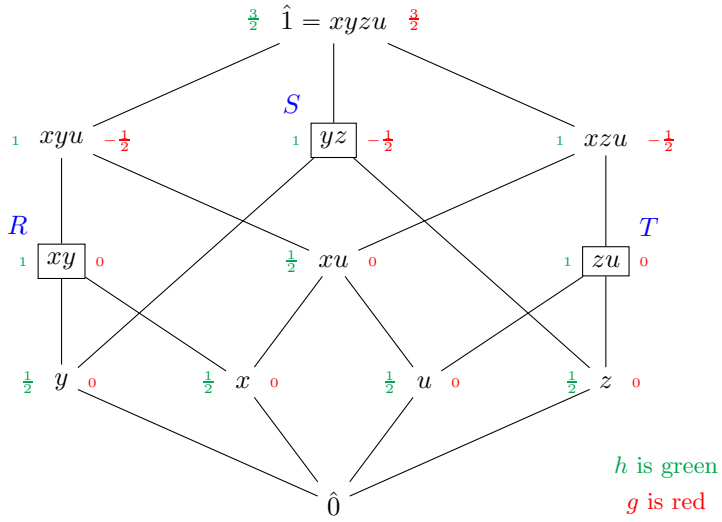

Figure 1: Q :- $R(x,y), S(y,z), T(z,u), xz \to u, yu \to x$

Lattice presentations and queries with functional dependencies are in 1-to-1 correspondence, up to the addition/removal of variables functionally equivalent to some other set of variables. To see this in one direction, consider any pair $(\mathbf{L}, \mathbf{R})$ where $\mathbf{R} = \{R_1, \ldots, R_m\} \subseteq L$ and $\bigvee \mathbf{R} = \hat{1}$. Let $\mathbf{X}$ be the join-irreducibles of $\mathbf{L}$, define a query (Q, FD) as follows. Let R_j be a relation with attributes Λ_{R_j}, and define FD such that the closed sets are $\{\Lambda_U \mid U \in L\}$ (in other words $\mathsf{FD} \stackrel{\text{def}}{=} \{X \to \Lambda_{\bigvee X} \mid X \subseteq \mathbf{X}\}$). One can check that the lattice presentation of this query is isomorphic to $(\mathbf{L}, \mathbf{R})$. In the other direction, consider a query with functional dependencies, (Q, FD). Call a variable x redundant if $Y \leftrightarrow x$ for some set Y that does not contain x. W.l.o.g. we can remove all redundant variables from Q (and decrease accordingly the arity of the relations containing x), because the values of a redundant variable can be recovered through expansion (Sec.2). We claim $x \mapsto x^+$ is a 1-to-1 mapping between the variables of Q and the join-irreducibles of $\mathbf{L}_{\mathsf{FD}}$. We first check that x^+ is join-irreducible: if $x^+ = Y \vee Z = (Y \cup Z)^+$ with $x \notin Y, x \notin Z$ then $x \leftrightarrow Y \cup Z$ contradicting the fact that x is not redundant. For injectivity, if $x \neq y$, then $x^+ = y^+$ implies $x \leftrightarrow y^+ - \{x\}$, again contradicting non-redundancy. Finally, surjectivity follows from the fact that $Y = \bigvee_{x \in Y} x^+$ for any closed set $Y \in L$: if Y is join-irreducible then $Y = x^+$ for some $x \in Y$, proving surjectivity.

Thus, in the rest of the paper we shall freely switch between queries and lattices, using the following correspondence:

variable $x \in \mathbf{X}$	join-irreducible $X \in L$
input $R_j(X_j)$	input $R_j \in L$ with variables Λ_{R_j}
set of variables $X \subseteq \mathbf{X}$	set of join-irreducibles X
its closure X^+	$X^+ \stackrel{\text{def}}{=} \Lambda_Z$, where $Z = \bigvee X$
FD= $\varnothing$	Boolean Algebra

The following simple observation illustrates the power of the lattice formalism.

Proposition 3.2. *If all functional dependencies are simple then* $\mathbf{L}$ *is a distributive lattice.*

Proof. Construct a directed graph $D = (V, E)$ where V is the set of all variables and there is an edge $(a, b) \in D$ if $a \to b$ is a functional dependency. Collapse each strongly connected component of this graph into a single vertex. We are left with a directed acyclic graph (DAG), which defines a poset P where $y \prec_P x$ iff there is a directed path from x to y in the DAG. Then, L is precisely the order ideal lattice of the direct sum of these posets. The claim follows because any order ideal lattice is distributive. □

It is easy to construct examples when non-simple FDs still produce a distributive lattice; for example, the lattice for $Q :- R(x), S(y), xy \to z$ is isomorphic to the Boolean algebra lattice $2^{x,y}$ which is distributive. Hence, the class of distributive lattices strictly covers the simple FD case.

Consider the non-distributive lattice M_3 (Fig. 3), then the procedure we discussed above associates it with the following query: $R(x), S(y), T(z), xy \to z, xz \to y, yz \to x$.

3.2 Database Instances for a Lattice

Let $\mathbf{X}$ denote the set of join-irreducibles of a lattice $\mathbf{L}$, and let $\mathsf{FD} = \left\{X \to \Lambda_{\bigvee X} \mid X \subseteq \mathbf{X}\right\}$ be all fd's implicitly defined by $\mathbf{L}$. A *database instance D for* $\mathbf{L}$ is a relation with attributes $\mathbf{X}$ that satisfies FD. We denote $h_D : L \to \mathbb{R}^+$ its entropy function, as defined in Sec. 2. If $(\mathbf{L}, \mathbf{R})$ is the lattice presentation of a query (Q, FD), then any database instance for $\mathbf{L}$ defines a standard database instance for Q by $R_j^D \stackrel{\text{def}}{=} \Pi_{R_j}(D)$. For example, given the query

$$Q :- R(x), S(y), T(z), xy \to z, xz \to y, yz \to x,$$

An example of an instance for its lattice (which is M_3 in Fig. 3) is:

$$D = \left\{(i, j, k) \in [N]^3 \mid i + j + k \mod N = 0\right\}.$$

It defines the following standard instance: $R^D = S^D = T^D = [N]$. Notice that all fd's are unguarded and hence lost in the standard instance. From Sec. 5.1 and up we will assume that during the query evaluation we have access to the UDF's that defined the unguarded fd's.

Definition 3.3 (Materializability)**.** A function $h : \mathbf{L} \to \mathbb{R}^+$ is called $\mathbf{L}$-*entropic* if $h = h_D$ for some instance D for $\mathbf{L}$ and it is called *materializable* w.r.t. an input query $(\mathbf{L}, \mathbf{R})$ if there exists an instance D such that $\log|R_j^D| \leqslant h(R_j)$ forall $j = 1, m$ and $\log|Q^D| \geqslant h(\hat{1})$. Obviously, any entropic function is materializable w.r.t. some input instance.

3.3 The Lattice Linear Program

This section shows that a very simple linear program defined on the FD lattice $\mathbf{L}$ is equivalent to GLVV bound. Given a query $(\mathbf{L}, \mathbf{R})$ and log cardinalities $(n_j)_{j=1}^m$, we define the following *Lattice Linear Program*, or LLP, over the *non-negative* variables $(h(X))_{X \in L}$:

$$\begin{array}{lrcll} \max & h(\hat{1}) & & & \\ & h(X \wedge Y) + h(X \vee Y) - h(X) - h(Y) & \leqslant & 0, & \forall X, Y \in L, X \not\leqslant Y. \\ & h(R_j) & \leqslant & n_j, & \forall j \in [m]. \end{array} \tag{5}$$

A feasible solution h to LLP is called a (non-negative) $\mathbf{L}$-*sub-modular* function. If h is also $\mathbf{L}$-*monotone*, i.e. $X \preceq Y$ implies $h(X) \leqslant h(Y)$, then h is called an $\mathbf{L}$-*polymatroid*. When $\mathbf{L}$ is clear from the context we drop the $\mathbf{L}$-prefix from $\mathbf{L}$-monotone, $\mathbf{L}$-submodular, and $\mathbf{L}$-polymatroid, respectively. We did not require h to be monotone because, at optimality, h^* can always be taken to be a polymatroid thanks to *Lovász's monotonization*: if h is non-negative $\mathbf{L}$-submodular, then the function $\bar{h}(\hat{0}) \stackrel{\text{def}}{=} 0$ and $\bar{h}(X) \stackrel{\text{def}}{=} \min_{Y : X \preceq Y} h(Y)$, $X \neq \hat{0}$, is an $\mathbf{L}$-polymatroid (see e.g. [21], pp. 774, and Appendix B) and satisfies: $\bar{h}(\hat{1}) = h(\hat{1})$ and $\forall X$, $\bar{h}(X) \leqslant h(X)$.

Proposition 3.4. *Let h^* be an optimal solution to the LLP of a query $(\mathbf{L}, \mathbf{R})$ where $N_j = |R_j|$, $j \in [m]$. Then, $2^{h^*(\hat{1})} = \mathsf{GLVV}(Q, \mathsf{FD}, (N_j)_j)$.*

Proof. From the discussion above, we can assume h^* is an $\mathbf{L}$-polymatroid. Define a function $h : 2^{\mathbf{X}} \to \mathbb{R}$ by $h(X) = h^*(X^+)$; then, clearly h is non-negative and monotone (on the original variable set, not on the

lattice $\mathbf{L}$), and satisfies the cardinality and fd constraints. We verify that it is also sub-modular: for any (not necessarily closed) set X of variables, noting that $(X \cap Y)^+ \preceq X^+ \wedge Y^+$, we have

$$\begin{aligned} h(X \cup Y) + h(X \cap Y) &= h^*((X \cup Y)^+) + h^*((X \cap Y)^+) \\ \text{(monotonicity)} &\leqslant h^*((X \cup Y)^+) + h^*(X^+ \wedge Y^+) \\ &= h^*(X^+ \vee Y^+) + h^*(X^+ \wedge Y^+) \\ \text{($\mathbf{L}$-submodularity)} &\leqslant h^*(X^+) + h^*(Y^+) \\ &= h(X) + h(Y). \end{aligned}$$

Conversely, given any polymatroid h that satisfies the cardinality and fd-constraints, the restriction of h on L is a feasible solution to LLP. □

The main goal in this paper is to design algorithms that compute a query Q in time $\tilde{O}(2^{h^*(\hat{1})})$. When the bound in the proposition is tight (which is currently an open problem) then such an algorithm is optimal, hence our second goal is to study cases when the bound is tight, and this happens if and only if the LLP has some optimal solution h^* that is materializable. A secondary goal in this paper is to find sufficient conditions for h to be materializable.

When $\mathsf{FD} = \varnothing$, $\mathbf{L}$ is a Boolean algebra. In this case it is easy to see that $\mathsf{AGM}(Q) = 2^{h^*(\hat{1})}$. To see this, in one direction start from an optimal weighted fractional vertex packing $(v_i^*)_{i=1}^k$ of the query hypergraph (Sec. 2), then the following is a feasible solution to LLP with the same objective value:

$$h(X) \stackrel{\text{def}}{=} \sum_{i: x_i \in X} v_i^* \tag{6}$$

Conversely, given a polymatroid h^* that is optimal to LLP, define $v_i \stackrel{\text{def}}{=} h^*(\{x_1, \ldots, x_i\}) - h^*(\{x_1, \ldots, x_{i-1}\})$, then it can be verified (using sub-modularity of h^*) that $(v_i)_{i=1}^k$ is a feasible weighted fractional vertex packing with the same objective value.

3.4 Embeddings

Embeddings allow us to construct an instance for one lattice $\mathbf{L}$ from an instance for a different lattice $\mathbf{L}'$.

Definition 3.5. An *embedding* between two lattices $\mathbf{L}$ and $\mathbf{L}'$ is a function $f : L \to L'$ such that $f(\bigvee X) = \bigvee f(X)$ forall $X \subseteq L$,[3] and $f(\hat{1}_{\mathbf{L}}) = \hat{1}_{\mathbf{L}'}$. An embedding between two queries $f : (\mathbf{L}, \mathbf{R}) \to (\mathbf{L}', \mathbf{R}')$ is an embedding from $\mathbf{L}$ to $\mathbf{L}'$ that is a bijection from $\mathbf{R}$ to $\mathbf{R}'$.

If f is an embedding from $\mathbf{L}$ to $\mathbf{L}'$ and h' is a non-negative, sub-modular function on $\mathbf{L}'$, then one can check that $h \stackrel{\text{def}}{=} h' \circ f$ is also sub-modular. If f is an embedding between queries, then their relation symbols are in 1-1 correspondence: $R_1, \ldots, R_m$ and $R'_1, \ldots, R'_m$ respectively. Fix two queries $(Q, \mathsf{FD}), (Q', \mathsf{FD}')$ with variables $\mathbf{X}, \mathbf{X}'$ respectively. Call a function $\mathcal{L} : 2^{\mathbf{X}} \to 2^{\mathbf{X}'}$ a *variable renaming* if $\mathcal{L}(X) = \mathcal{L}(X^+)$, $\mathcal{L}(X) = (\bigcup_{x \in X} \mathcal{L}(x))^+$ ($\mathcal{L}$ is uniquely defined by its values on single variables), and $\mathcal{L}(\mathbf{X}) = \mathbf{X}'$ (all variables in Q' are used). One can check that every embedding f defines the variable renaming $\mathcal{L}(X) \stackrel{\text{def}}{=} f(X^+)$ forall $X \subseteq \mathbf{X}$, and vice versa.

Fix an embedding $f : \mathbf{L} \to \mathbf{L}'$, and an instance D' for $\mathbf{L}'$. Define the database instance $D \stackrel{\text{def}}{=} f^{-1}(D')$ as $D \stackrel{\text{def}}{=} \{t \circ \mathcal{L} \mid t \in D'\}$. In other words, every tuple $t' \in D'$ has attributes $\mathbf{X}'$, hence it can be viewed as a function $t' : \mathbf{X}' \to \text{Dom}$: for each such t' we include in D a tuple $t' \circ \mathcal{L} : \mathbf{X} \to \text{Dom}$, obtained from t' by renaming its attributes. One can check that $|D| = |D'|$ (because all variables in $\mathbf{X}'$ are used), and the same holds for all projections $|\Pi_X(D)| = |\Pi_{\mathcal{L}(X)}(D')|$ forall $X \subseteq \mathbf{X}$. For example, if $\mathcal{L}(x_1) \stackrel{\text{def}}{=} \{y_1, y_2\}$, $\mathcal{L}(x_2) \stackrel{\text{def}}{=} \{y_2, y_3\}$ then $D(x_1, x_2) \stackrel{\text{def}}{=} \{((ab), (bc)) \mid (a, b, c) \in D'(y_1, y_2, y_3)\}$. This proves:

[3] Equivalently: f is the left adjoint of a Galois connection.

Proposition 3.6. *If $f : \mathbf{L} \to \mathbf{L}'$ is a lattice embedding, then for any instance D' for $\mathbf{L}'$, $h_D = h_{D'} \circ f$, where $D \stackrel{\text{def}}{=} f^{-1}(D')$. If $f : (\mathbf{L}, \mathbf{R}) \to (\mathbf{L}', \mathbf{R}')$ is a query embedding , and $h' : L' \to \mathbb{R}^+$ is materializable, then $h \stackrel{\text{def}}{=} h' \circ f$ is also materializable.*

Definition 3.7. A *quasi-product database instance* for $\mathbf{L}$ is an instance of the form $f^{-1}(D')$ where f embeds $\mathbf{L}$ into a Boolean algebra A, and D' is a product database instance for A.

Example 3.8. Continuing our running example, consider the following two queries:

$$\begin{aligned} Q &= R(x,y), S(y,z), T(z,u), xz \to u, yu \to x \\ Q' &= R'(a,b), S'(b,c), T'(c,a) \end{aligned}$$

Q is the query in Fig. 1, while Q' is the triangle query (without fd's). Consider the following renaming $\mathcal{L}(x) = \mathcal{L}(u) = a$, $\mathcal{L}(y) = b$, $\mathcal{L}(z) = c$, which defines an embedding[4] from the lattice in Fig. 1 to the Boolean algebra $2^{\{a,b,c\}}$. From the product instance

$$D' \stackrel{\text{def}}{=} \left\{ (a,b,c) = (i,j,k) \mid i,j,k \in [\sqrt{N}] \right\}$$

we construct, through renaming, the quasi-product instance

$$D = \left\{ (x,y,z,u) = (i,j,k,i) \mid i,j,k \in [\sqrt{N}] \right\}.$$

Notice that the FD's $xz \to u, yu \to x$ hold in D (as they should).

3.5 Inequalities

In information theory, an *information inequality* is defined by a vector of real numbers $(w_X)_{X \subseteq \mathbf{X}}$ such that the inequality $\sum_{X \subseteq \mathbf{X}} w_X H(X) \geqslant 0$ holds for any entropic function H. We introduce here a related notion, which we use to describe upper bounds on the query size.

Fix a query $(\mathbf{L}, \mathbf{R})$. An *output-inequality* is given by a vector of non-negative real numbers $(w_j)_{j \in [m]}$ such that the inequality

$$\sum_{j=1}^{m} w_j h(R_j) \geqslant h(\hat{1}) \tag{7}$$

holds for all polymatroids h on the lattice $\mathbf{L}$. For example, in a Boolean algebra output inequalities correspond to Shearer's lemma: $\sum_j w_j h(R_j) \geqslant h(\mathbf{X})$ iff $(w_j)_{j=1}^m$ is a fractional edge cover of the hypergraph with nodes $\{x \mid x \in \mathbf{X}\}$ and hyperedges $\{R_j \mid j \in [m]\}$. *Any* output inequality gives us immediately an upper bound on the output size of a query, because (7) implies $h(\hat{1}) \leqslant \sum_j w_j n_j$ for any feasible solution h of the LLP with log-cardinalities $(n_j)_{j=1}^m$.

An ouput size upperbound is best if it is minimized, which is precisely the objective of the dual-LLP, which is defined over *non-negative* variables $(w_j)_{j=1}^m$ and $(s_{X,Y})_{X \not\leqslant Y}$, corresponding to the input R_j and incomparable pairs of lattice elements (thus $s_{X,Y}$ is the same variable as $s_{Y,X}$):

$$\begin{aligned} \min \quad & \sum_j w_j n_j \\ & \sum_{X \not\leqslant Y : X \vee Y = \hat{1}} s_{X,Y} \geqslant 1, \\ & w_j + \sum_{X \not\leqslant Y : X \vee Y = R_j} s_{X,Y} + \sum_{X \not\leqslant Y : X \wedge Y = R_j} s_{X,Y} - \sum_{X \not\leqslant R_j} s_{X,R_j} \geqslant 0, \quad \forall j \in [m] \end{aligned} \tag{8}$$

[4]We need to check $\mathcal{L}(X) = \mathcal{L}(X^+)$ for all X: e.g. $\mathcal{L}(xz) = \mathcal{L}(xz^+)$ holds because $\mathcal{L}(xz) = ac$, $\mathcal{L}(xz^+) = \mathcal{L}(xzu) = ac$, and similarly $\mathcal{L}(yu) = \mathcal{L}(yu^+)$.

A simple way to prove an output inequality is to write it as a non-negative linear combination of sub-modularity inequalities. (See example below.) The following Lemma shows that every output inequality can be proven this way, as follows from the theory of generalized inequalities [6, Ch.2].

Lemma 3.9. *Let M denote the matrix of sub-modular inequalities (first inequality of the LLP (5)). Given a non-negative vector $(w_j)_{j\in[m]}$, define the vector $(c_X)_{X\in L}$ by setting $c_{\hat{1}} \stackrel{\text{def}}{=} 1$, $c_{R_j} \stackrel{\text{def}}{=} -w_j$ and $c_X \stackrel{\text{def}}{=} 0$ otherwise. Then, the following statements are equivalent*

(i) (7) *holds for all polymatroids h*

(ii) (7) *holds for all non-negative sub-modular functions h*

(iii) there exists $s \geqslant 0$ such that $c^T \leqslant s^T M$. Equivalently, w is part of a feasible solution (s, w) of the dual-LLP (8).

Furthermore, if (s^*, w^*) and h^* are dual- and primal-optimal solutions, then $h^*(\hat{1}) = \sum_{j=1}^m w_j^* n_j$.

Proof. Using the Lovász monotonization map $h \to \bar{h}$ described in Sec 3.3, it is straightforward to show that (i) and (ii) are equivalent. To see the equivalence between (ii) and (iii), note that $K \stackrel{\text{def}}{=} \{h \mid Mh \leqslant 0, h \geqslant 0\}$ is the set of all non-negative, sub-modular functions, and the polar cone $K^* \stackrel{\text{def}}{=} \{c \mid c^T h \leqslant 0, \forall h \in K\}$ is the set of all inequalities that hold for all non-negative sub-modular functions. Now, $c \in K^*$ iff the LP $\max_h\{c^T h \mid Mh \leqslant 0, h \geqslant 0\}$ has objective value 0, which holds iff the dual LP $\min_s \{0 \mid c^T \leqslant s^T M, s \geqslant 0\}$ is feasible. The last statement is strong duality. □

Example 3.10. The lattice presentation of the triangle query $Q = R(x, y), S(y, z), T(z, x)$ (without FD's) is the Boolean algebra $2^{\{x,y,z\}}$. The following output inequality:

$$h(xy) + h(yz) + h(zx) \quad \geqslant \quad 2h(\hat{1}) \tag{9}$$

follows by adding two sub-modularity inequalities $h(xy) + h(yz) \geqslant h(\hat{1}) + h(y)$ and $h(y) + h(zx) \geqslant h(\hat{1})$. It corresponds to the dual solution $s_{xy,yz} = s_{y,zx} = 1$ with the rest 0. Some other output inequalities are $h(xy) + h(xz) \geqslant h(\hat{1})$, $h(xy) + h(yz) \geqslant h(\hat{1})$, and $h(xz) + h(yz) \geqslant h(\hat{1})$. Together, these four output inequalities prove Eq.(4).

4 Normal lattices

We show here that an $\mathbf{L}$-polymatroid h can be materialized as a quasi-product database instance iff it satisfies all output inequalities given by fractional edge covers of a certain hypergraph; in that case we call h *normal*. The normal polymatroids are the largest class of polymatroids that preserve the elegant properties of the AGM bound: upper bound given in terms of a fractional edge cover, and lower bound given a (quasi-) product database. We then extend normality to a lattice, which we call normal if its optimal polymatroid h^* is normal.

Recall the Möbius inversion formula in a lattice $\mathbf{L}$:

$$h(X) = \sum_{Y : X \preceq Y} g(Y) \quad \text{iff} \quad g(X) = \sum_{Y : X \preceq Y} \mu(X, Y) h(Y) \tag{10}$$

where $\mu(X, Y)$ is the Möbius function on $\mathbf{L}$ [22]. In information theory, when h is an entropy and $\mathbf{L}$ a Boolean algebra, the quantity $-g(X)$ is the (multivariate) *conditional mutual information* $I(\hat{1} - X \mid X)$, which we abbreviate CMI. For example, in the Boolean Algebra $2^{\{x,y,z\}}$:

$$\begin{aligned}
g(xyz) &= h(xyz) \\
g(xy) &= h(xy) - h(xyz) \\
g(x) &= h(x) - h(xy) - h(xz) + h(xyz) \\
g(\hat{0}) &= h(\hat{0}) - h(x) - h(y) - h(z) + h(xy) + h(xz) + h(yz) - h(xyz)
\end{aligned}$$

We give below a simple sufficient condition on g which implies that h, defined by Eq.(10) is a polymatroid. We need:

Lemma 4.1. *For $\mathcal{S} \subseteq L$, let $c(Z, \mathcal{S}) \stackrel{\text{def}}{=} |\{U \mid U \in \mathcal{S}, U \preceq Z\}|$. Then, for all $X, Y, Z \in L$, $c(Z, \{X, Y\}) \leqslant c(Z, \{X \wedge Y, X \vee Y\})$; moreover, if $\mathbf{L}$ is a distributive lattice and Z is a meet-irreducible, then equality holds.*

Proof. If $c(Z, \{X, Y\}) = 1$, then w.l.o.g. we assume $X \preceq Z$, hence $X \wedge Y \preceq Z$, and $c(Z, \{X \wedge Y, X \vee Y\}) \geqslant 1$. If $c(Z, \{X, Y\}) = 2$, then both $X \preceq Z, Y \preceq Z$ hold, hence both $X \wedge Y \preceq Z$, $X \vee Y \preceq Z$ hold, hence $c(Z, \{X \wedge Y, X \vee Y\}) = 2$. Next, suppose $\mathbf{L}$ is distributive and Z is a meet-irreducible. If $c(Z, \{X \wedge Y, X \vee Y\}) = 2$ then $X \vee Y \preceq Z$ hence both $X \preceq Z, Y \preceq Z$ hold. If $c(Z, \{X \wedge Y, X \vee Y\}) = 1$, then $X \wedge Y \preceq Z$. By distributivity $Z \succeq (X \wedge Y) \vee Z = (X \vee Z) \wedge (Y \vee Z)$. On the other hand both $Z \preceq X \vee Z$ and $Z \preceq Y \vee Z$ hence $Z = (X \vee Z) \wedge (Y \vee Z)$. Since Z is a meet-irreducible, it must be equal to one of the two terms. Assume w.l.o.g. $Z = X \vee Z$ then $X \preceq Z$ proving $c(Z, \{X, Y\}) \geqslant 1$. □

Lemma 4.1 immediately implies:

Lemma 4.2. *Let g be any function s.t. $g(Z) \leqslant 0$ for $Z \prec \hat{1}$, and $g(\hat{1}) = -\sum_{Z \prec \hat{1}} g(Z)$. Then the function h defined by Eq.(10) is a polymatroid. Furthermore, if $\mathbf{L}$ is distributive, then h is a modular polymatroid, i.e. $h(X) + h(Y) = h(X \vee Y) + h(X \wedge Y)$ for all $X, Y \in L$.*

Proof. Non-negativity and monotonicity are easy to verify, submodularity follows from:

$$\begin{aligned} h(X) + h(Y) &= \sum_{Z: X \preceq Z} g(Z) + \sum_{Z: Y \preceq Z} g(Z) \\ &= \sum_{Z} c(Z, \{X, Y\}) \cdot g(Z) \\ &\geqslant \sum_{Z} c(Z, \{X \vee Y, X \wedge Y\}) \cdot g(Z) \\ &= h(X \vee Y) + h(X \wedge Y), \end{aligned}$$

where the inequality holds due to Lemma 4.1 and the fact that $g(Z) \leqslant 0$ for $Z \neq \hat{1}$. When $\mathbf{L}$ is distributive, the inequality becomes an equality. □

Any function h satisfying the property stated in Lemma 4.2 is called a *normal submodular function*. If, furthermore, $g(Z) = 0$ for all $Z \prec \hat{1}$ other than the co-atoms, then we say that h is *strictly normal*.

For a simple example, consider the function h in our running example Fig.1: it is strictly normal because it is defined by the CMI g shown in the figure. A negative example is given by function h on the left of Fig. 3, which is not normal, because its CMI satisfies $g(\hat{0}) > 0$. If the lattice $\mathbf{L}$ is a Boolean algebra, then the optimal polymatroid h^* given given by Eq.(6) is strictly normal: its CMI is $g(\mathbf{X}) = \sum_i v_i^*$, $g(\mathbf{X} - \{x_i\}) = -v_i^*$, and $g = 0$ everywhere else.

Normal polymatroids are precisely non-negative linear combinations of "step functions". For every $Z \in L$, *the step function h_Z at Z* is defined by $h_Z(X) = 1$ if $X \not\preceq Z$, and 0 otherwise. Every step function is normal, because its Möbius inverse is $g_Z(\hat{1}) = 1$, $g_Z(Z) = -1$, and $g_Z(X) = 0$ otherwise. Any non-negative linear combination of step functions is normal. Conversely, if h is a normal polymatroid, then denoting $a_Z = -g(Z)$ for all $Z \neq \hat{1}$ we have $a_Z \geqslant 0$ (since h is normal) and $g = \sum_Z a_Z g_Z$, implying $h = \sum_Z a_Z h_Z$.

4.1 Connection to Quasi-Product Instance

We have seen that the worst-case instance of the AGM-bound is a product database instance (Theorem 2.1), and its entropy function given in Eq.(6) is strictly normal. We generalize this observation by proving that normal polymatroids are precisely entropy functions of quasi-product instances (Definition 3.7). For one direction we need:

Lemma 4.3. *Let $f : L \to L'$ be an embedding between two lattices, and let h' be a normal polymatroid on $\mathbf{L}'$. Then $h \stackrel{\text{def}}{=} h' \circ f$ is a normal polymatroid on $\mathbf{L}$.*

Proof. Recall that an embedding is the left adjoint of a Galois connection. Let $r : L' \to L$ be its right adjoint[5], in other words $f(X) \preceq Y$ iff $X \preceq r(Y)$. Let g' be the CMI for h'. The function $g(X) \stackrel{\text{def}}{=} \sum_{Y:r(Y)=X} g'(Y)$ is the CMI of h, because, forall $X \in L$:

$$\sum_{Z:X\preceq Z} g(Z) = \sum_{Z:X\preceq Z} \sum_{Z':r(Z')=Z} g'(Z') = \sum_{Z':X\preceq r(Z')} g'(Z') = \sum_{Z':f(X)\preceq Z'} g'(Z') = h'(f(X)) = h(X)$$

In any Galois connection, $r(\hat{1}) = \hat{1}$ (since $f(\hat{1}) \preceq \hat{1}$ iff $\hat{1} \preceq r(\hat{1})$). By normality of h' we have $Y \neq \hat{1}$ implies $g'(Y) \leqslant 0$. Therefore, $X \neq \hat{1}$ implies $g(X) = \sum_{Y:r(Y)=X} g'(Y) \leqslant 0$, proving that h is normal. □

In the opposite direction, we need:

Definition 4.4. Let $A = 2^{\mathbf{X}}$ be a Boolean algebra. The canonical instance is the product database $D = [2]^{\mathbf{X}}$; note that $|\Pi_X(D)| = 2^{|X|}$, and thus $h_D(X) = |X|$, forall $X \subseteq \mathbf{X}$.

Let $\mathbf{L}$ be a lattice and h an integer-valued, normal polymatroid, and g its CMI. For all $X \in L$, $X \neq \hat{1}$, let $C(X)$ be a set of $-g(X)$ arbitrary elements, such that the sets $(C(X))_{X\neq\hat{1}}$ are disjoint. Define $C(\hat{1}) = \varnothing$, and $\mathbf{C} \stackrel{\text{def}}{=} \bigcup_X C(X)$. The *canonical Boolean Algebra* is $A = (2^{\mathbf{C}}, \supseteq)$, and the *canonical embedding* of $\mathbf{L}$ is $f : L \to A$, $f(X) \stackrel{\text{def}}{=} \bigcup_{Z:X\preceq Z} C(Z)$.

A is an "upside-down" Boolean algebra, where $Y \preceq Y'$ iff $Y \supseteq Y'$, and $\vee$ is set intersection; one can check that f commutes with $\vee$. We can now prove our result. (As in the AGM bound, there is a bit of loss when h are not integral, but this does affect asymptotics)

Lemma 4.5. *Let h be an integral, non-negative, submodular function h on a lattice $\mathbf{L}$. Then h is normal iff it is the entropy function of a quasi-product instance.*

Proof. In one direction, let D is a quasi-product instance and h_D be its entropy function. By definition, $D = f^{-1}(D')$, where $f : L \to A$ is an embedding into a Boolean algebra and D' is a product instance for A. Since $h_{D'}$ is normal (even strictly normal), h_D is also normal by Lemma 4.3. In the opposite direction, assume h is normal. Let $f : L \to A$ be the canonical embedding, and D' be the canonical instance for A. Then $D \stackrel{\text{def}}{=} f^{-1}(D')$ is the quasi-product instance, and one can check that $h = h_D = h_{D'} \circ f$. □

Example 4.6. We illustrate the construction above on the lattice $\mathbf{L}$ in Fig. 1, showing that it leads to the quasi-product instance in Example 3.8. The canonical embedding is into the Boolean algebra $A = 2^{\{a,b,c\}}$ because there are three co-atoms in $\mathbf{L}$: if we call the *atoms* of this Boolean algebra a, b, c, then the embedding $f : L \to 2^{\{a,b,c\}}$ is precisely the renaming in Example 3.8. The polymatroid h is equal to $h' \circ f$, where h' is defined on A by $h'(a) = h'(b) = h'(c) = 1/2$, $h'(ab) = h'(ac) = h'(bc) = 1$, $h'(abc) = 3/2$, whose materialization is that in Example 3.8 for $N = 2$.

4.2 Connection to Fractional Edge Covering

Definition 4.7. Let $(\mathbf{L}, \mathbf{R})$ be a query given in lattice presentation. The co-atomic hypergraph $H_{\text{co}} = (V_{\text{co}}, E_{\text{co}})$ is defined as follows. The nodes V_{co} are the co-atoms of $\mathbf{L}$, and $E_{\text{co}} = \{e_1, \ldots, e_m\}$ with $e_j = \{Z \mid Z \in V_{\text{co}}, R_j \not\preceq Z, R_j \in \mathbf{R}\}$. In other words, each relation R_j defines a hyperedge e_j consisting of those nodes that do *not* contain the variables of R_j. A simple illustration of a co-atomic hypergraph is in Fig.2.

Lemma 4.8. *Inequality (7) holds for all normal polymatroids h iff it holds for all strictly normal polymatroids iff $(w_j)_{j=1}^m$ is a fractional edge cover of H_{co}.*

[5]The standard notation is g, but we use g for the CMI.

Proof. Since each (stricly) normal polymatroid is a non-negative linear combination of (co-atomic) step functions, it suffices to assume that h is a step function. First consider a co-atomic step functions h_Z, i.e. Z is a co-atom: the left hand size of Eq. (7) is $\sum_{j:Z\in e_j} w_j$, and the right hand side is 1, meaning that the inequality holds iff $(w_j)_{j=1}^m$ covers node $Z \in V_{\text{co}}$. It remains to show that if (7) holds for all co-atomic step functions then it holds for all step functions h_X: let Z be any a co-atom s.t. $Z \succeq X$, then $h_Z \leqslant h_X$ and $h_Z(\hat{1}) = h_X(\hat{1}) = 1$ implying that (7) holds for h_X. □

The co-atomic hypergraph is the natural concept to capture Shearer's lemma in a general lattice. In fact, every vector $(w_j)_{j=1}^m$ for which the output inequality (7) holds is a fractional edge cover of the co-atomic hypergraph. One may wonder whether the output inequalities could also be described by the *atomic hypergraph*, defined in a similar way. In a Boolean algebra $2^{\mathbf{X}}$ the atomic and co-atomic hypergraphs are isomorphic via $x \mapsto \mathbf{X} - \{x\}$, since $x \in R_j$ iff $R_j \nsubseteq \mathbf{X} - \{x\}$, but in a general lattice the atomic hypegraph does not seem to lead to any interesting properties.

4.3 Normal Lattices

A normal lattice is a lattice where, at optimality, the polymatroid is normal. This is captured by the following, which is the main result of this section.

Theorem 4.9. *Let $(\mathbf{L}, \mathbf{R})$ be a query in lattice presentation. The following are equivalent:*

1. *For every non-negative, submodular function h there exists a normal polymatroid h' s.t. $h'(\hat{1}) = h(\hat{1})$ and $h'(R_j) \leqslant h(R_j)$ for all $R_j \in \mathbf{R}$.*
2. *The previous property holds and h' is strictly normal.*
3. *Inequality (7) holds for all non-negative, submodular functions h iff $w_1, \ldots, w_m$ is a fractional edge cover of the co-atomic hypergraph.*
4. *Every non-negative submodular function h on $\mathbf{L}$ has a materialization that is a quasi-product database instance.*

If L satisfies any of these conditions, then we call it a normal lattice w.r.t. $\mathbf{R}$*. If L is normal w.r.t. any inputs $\mathbf{R}$, then we call it shortly a* normal lattice*.*

Proof. We first prove that items 1, 2, and 3 are equivalent, then show that item 2 is equivalent to 4.

Item 2⇒ item 1 is obvious.

For Item 1⇒ item 3, let h be any submodular function on L. Let h' be a normal function as defined in item 1. Since h' is normal, it satisfies the inequality (7); it suffices to note that $\sum_j w_j h(R_j) \geqslant \sum_j w_j h'(R_j) \geqslant h'(\hat{1}) \geqslant h(\hat{1})$.

For Item 3⇒ item 2, let h be any submodular function on L, and let $Z_1, \ldots, Z_k$ be all co-atoms in L. To define a strictly normal function h', we need to find k numbers $a_i \geqslant 0$, $i \in [k]$ and define:

$$\begin{aligned} h'(\hat{1}) &= \sum_i a_i \\ h'(X) &= h'(\hat{1}) - \sum \{a_i \mid i \in [k] : X \preceq Z_i\} = \sum \{a_i \mid i \in [k] : X \npreceq Z_i\} \end{aligned}$$

We need to find these numbers such that $h'(\hat{1}) \geqslant h(\hat{1})$ and $h'(R_j) \leqslant h(R_j)$ for all $R_j \in \mathbf{R}$. To do that, consider the following linear program:

$$\begin{array}{ll} \text{maximize} & a_1 + \ldots + a_k \\ \forall j \in [m]: & \sum \{a_i \mid i \in [k], R_j \npreceq Z_i\} \leqslant h(R_j) \end{array}$$

We claim that its optimal value is $\geqslant h(\hat{1})$: this implies that h' defined above is a strictly normal function satisfying the requirement in the theorem. To prove the claim, consider the dual LP. Its variables are b_j, for $j \in [m]$:

$$\begin{array}{ll} \text{minimize} & \sum_j b_j h(R_j) \\ \forall i \in [k]: & \sum \{b_j \mid j \in [m], R_j \not\perp Z_i\} \geqslant 1 \end{array}$$

Thus, the feasible solutions of the dual LP are precisely the fractional edge covers of the dual hypergraph. Hence, by assumption in Item 3, the output inequality (7) holds:

$$\sum_j b_j h(R_j) \geqslant h(\hat{1})$$

which completes the proof.

Finally, we show that item 2 holds iff 4 holds. For the only if direction, consider a non-negative, submodular function h; by 1 there exists a normal polymatroid h' s.t. $h'(\hat{1}) = h(\hat{1})$ and $h'(R_j) \leqslant h(R_j)$ forall j. By Lemma 4.5, h' is the entropy of some quasi-product instance D. Then D is a materialization of h because $\log|R_j^D| = h'(R_j) \leqslant h(R_j)$ and $\log|D| = h'(\hat{1}) = h(\hat{1})$. Conversely, assume item 4 holds: h has some materialization D that is a quasi-product. Then by Lemma 4.5, h_D is normal, and $h_D(R_j) = \log|R_j^D| \leqslant h(R_j)$, $h_D(\hat{1}) = \log|D| = h(\hat{1})$, proving that item 1 holds, with $h' \stackrel{\text{def}}{=} h_D$. □

Normal lattices appear to be a new concept. It is decidable whether a lattice $\mathbf{L}$ is normal w.r.t. $\mathbf{R}$, using the following naive procedure. Enumerate all vertices $(w_j)_j$ of the fractional edge packing polytope of the co-atomic hypergraph, and check that the output inequality (7) holds, by using the criterion in Lemma 3.9.

We prove in Sec. 5.2 that every distributive lattice is normal; furthermore, Proposition 3.2 says that any set of unary FDs generate a distributive lattice, which is therefore normal. The lattice on the left of Fig.3 is a Boolean algebra, hence it is normal: note that h in the figure is not normal[6], but we can simply increase $h(\hat{1})$ to 3 and now it is normal.

The lattice $\mathbf{L}$ in Fig.1 is normal w.r.t. inputs xy, yz, zu, which follows by exhaustively proving all inequalities defined by the fractional edge covers of the co-atomic hypergraph, shown in Fig. 2. For example, the edge cover $(1/2, 1/2, 1/2)$ corresponds to the inequality $h(xy) + h(yz) + h(zu) \geqslant 2h(\hat{1})$, which can be proven by $h(xy) + h(yz) \geqslant h(\hat{1}) + h(y)$ and $h(y) + h(zu) \geqslant h(\hat{1}) + h(0)$. In fact, $\mathbf{L}$ is normal w.r.t. any inputs. Notice that $\mathbf{L}$ is not distributive.

The lattice M_3 on the right of Fig.3 is not normal. Its co-atomic hypergraph has edges $e_x = \{y, z\}$, $e_y = \{x, z\}$, $e_z = \{x, y\}$ and the inequality corresponding to the fractional edge cover $(1/2, 1/2, 1/2)$, $h(x) + h(y) + h(z) \geqslant 2h(\hat{1})$, fails for the submodular function shown in the figure. In particular, this polymatroid h is not materializable as a quasi-product database instance, but can be materialized as $\{(i, j, k) \mid i, j, k \in \{0, 1\}, i + j + k = 0 \mod 2\}$.

We give a necessary condition for normality (and we conjecture it is also sufficient):

Proposition 4.10. *Let L be a lattice that contains a sublattice $\{U, X, Y, Z, \hat{1}\}$ isomorphic to M_3, and let $\mathbf{R} = \{X, Y, Z\}$. Then L is not normal.*

Proof. The co-atomic hypergraph has three hyper-edges, X, Y, Z. Then $(1/2, 1/2, 1/2)$ is a fractional edge cover of the co-atomic hypergraph, because every co-atom is above at most one of X, Y, Z, because $X \vee Y = X \vee Z = Y \vee Z = \hat{1}$, hence it belongs to at least two hyper-edges of the hypergraph. This defines the inequality:

$$h(X) + h(Y) + h(Z) \geqslant 2h(\hat{1}).$$

[6]The function XOR on three variables, $R(x, y, z) = \{(a, b, c) \mid a, b, c \in \{0, 1\}, a \text{ xor } b \text{ xor } c = 0\}$, is the canonical example of a distribution whose entropy has a negative mutual information.

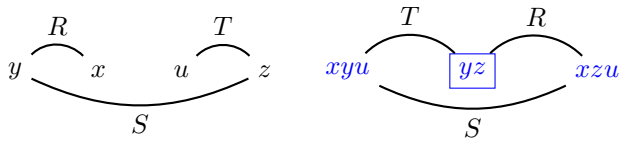

Atomic Hypergraph H_A Co-atomic Hypergraph H_C

Figure 2: Atomic and Co-atomic hypergraphs for Fig.1

We construct a polymatroid h that violates this inequality:

$$h(W) = \begin{cases} 0 & \text{when } W \preceq U \\ 1 & \text{when } W \preceq X \text{ or } W \preceq Y \text{ or } W \preceq z \\ 2 & \text{when } W \text{ is not below } X \text{ or } Y \text{ or } Z \end{cases}$$

Then $h(X) = h(Y) = h(Z) = 1$, $h(\hat{1}) = 2$ and therefore h violates the inequality above. We prove that h is a polymatroid. Monotonicity is easy to check. We prove that it satsifies the submodularity laws:

$$h(A) + h(B) \geqslant h(A \vee B) + h(A \wedge B)$$

Case 1: $h(A \vee B) = 0$. Then all elements are below U, and all h are 0.

Case 2: $h(A \vee B) = 1$. If $h(A) = 1$ then the inequality follows from $h(B) \geqslant h(A \wedge B)$, and similarly for $h(B) = 1$. If both $h(A) = h(B) = 0$ then both A, B are below U, hence $A \vee B$ is below U, contradicting $h(A \vee B) = 1$

Case 3: $h(A \vee B) = 2$. We can rule out the cases when $h(A), h(B)$ are 0,0 or 0,1 because that implies both A, B are below X (or Y or Z), implying $h(A \vee B) = 1$. Also, the case when $h(A) = 2$ or $h(B) = 2$ follows immediately from monotonicity. Thus, assume $h(A) = h(B) = h(A \wedge B) = 1$. Here we use the structure of M_3. If A and B are below two distinct elements from X, Y, Z, e.g. $A \leqslant X$ and $B \leqslant Y$, then $A \wedge B \leqslant U$ contradicting $h(A \wedge B) = 1$. Hence both A, B are below X. But that implies $A \vee B$ is also below X, contradicting $h(A \vee B) = 2$. □

Gottlob et al. [11] define a *coloring* of a query Q with variables $\mathbf{X}$ to be a function $\mathcal{L} : \mathbf{X} \to 2^{\mathbf{X}'}$ such that $\mathcal{L}(\mathbf{X}) \neq \varnothing$, and for any FD $X \to Y$, $\mathcal{L}(Y) \subseteq \mathcal{L}(X)$, where, for each set X, $\mathcal{L}(X) \stackrel{\text{def}}{=} \bigcup_{x \in X} \mathcal{L}(X)$. Then the *color number of* $\mathcal{L}$ is defined as $C(\mathcal{L}) = |\mathcal{L}(\mathbf{X})| / \max_j |\mathcal{L}(R_j)|$ and the *color number of* Q is $\max_{\mathcal{L}} C(\mathcal{L})$. They prove two results: if the functional dependencies are restricted to simple keys then $|Q^D| \leqslant \left(\max_j |R_j^D|\right)^{C(Q)}$, and moreover this bound is essentially tight, even for general functional dependencies.

Colorings correspond one-to-one to integral, normal polymatroids, via $h(x) = |\bigcup_{x \in X} \mathcal{L}(x)|$. To see this, in one direction let $\mathcal{L} : \mathbf{X} \to 2^{\mathbf{X}'}$ be a coloring and define $h(X) \stackrel{\text{def}}{=} |\bigcup_{x \in X} \mathcal{L}(x)|$. The function $f(X) \stackrel{\text{def}}{=} \bigcup_{x \in X} \mathcal{L}(x)$ is an embedding $f : \mathbf{L} \to 2^{\mathbf{X}'}$ into a Boolean algebra (we assume w.l.o.g. that $f(\mathbf{X}) = \mathbf{X}'$, otherwise we redefine $\mathbf{X}'$), and we have $h = h' \circ f$ where $h'(Y) \stackrel{\text{def}}{=} |Y|$: since h' is a strictly normal polymatroid Lemma 4.3 says that h is also normal. In the other direction, if h is an integral, normal polymatroid, then its canonical embedding (Definition 4.4) defines a coloring $\mathcal{L}$ s.t. $h(X) \stackrel{\text{def}}{=} |\bigcup_{x \in X} \mathcal{L}(x)|$. Thus, colorings are essentially normal polymatroids. This implies the two results in [11] as follows. If FD consists only of simple keys, then the lattice $\mathbf{L}$ is distributive (Prop. 3.2), hence it is normal (Corollary 5.23), hence at optimality h^* can be assumed to be a normal polymatroid (Theorem 4.9 item 1), equivalently a coloring. The second result, tightness, follows from our Lemma 4.5, since any normal polymatroid (coloring) is the entropy of a quasi-product database instance.

To summarize, we have introduced a framework for studying the query upper bound under FDs that uses the lattice of its closed sets of attributes, which extends and generalizes the framework introduced by Gottlob et al. [11].

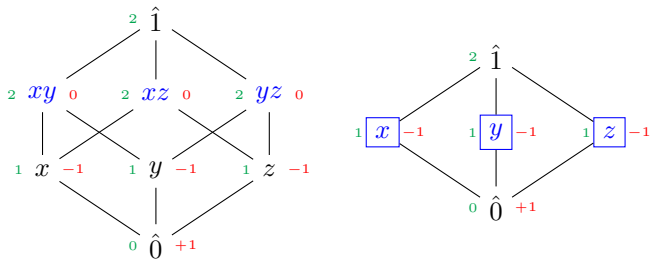

Figure 3: A non-normal function h (left) and the non-normal lattice M_3 (right)

5 Proof Sequence to Algorithm

We aim to design an algorithm that, given an input query Q and database instance D, computes Q^D in time $\tilde{O}(2^{h^*(\hat{1})})$, where h^* is the optimal solution to the LLP (5). From Lemma 3.9, if (s^*, w^*) is dual-optimal then $h^*(\hat{1}) = \sum_{j=1}^m w_j^* n_j$ and the output inequality (7) holds with $w_j = w_j^*$, for *all* polymatroids. Our main algorithmic theme is to turn this process "inside-out": find a "proof" of inequality (7) consisting of a sequence of steps transforming the left-hand-side into the right-hand-side, then interpret these steps as algorithmic steps whose runtime does not exceed what the symbolic step allows.

We show that three different proof techniques for (7) lead to three different bounds and algorithms. The first two are known techniques for proving Shearer's lemma (i.e. the Boolean algebra case) that we adapt to lattices. The bounds are tight and corresponding algorithms are worst-case optimal for some classes of lattices (such as distributive lattices, which include the Boolean algebra and the lattice of queries with simpld fd's, subsuming results from [11, 18, 23]). However, for a generic lattice the bounds these techniques can prove are not $h(\hat{1}) \leqslant \sum_j w_j^* h(R_j)$, and thus the corresponding algorithms' runtimes in general can be worse than $\tilde{O}(2^{h^*(\hat{1})})$. An advantage of these two algorithms is that there is only one log-factor hidden in the $\tilde{O}$.

The last proof technique is our own, designed specifically to achieve the optimal LLP bound in an arbitrary lattice; this leads to an algorithm called CSMA runing in the stated time of $\tilde{O}(2^{h^*(\hat{1})})$, where $\tilde{O}$ hides a polylogarithmic factor. This algorithms needs to regularize the data, somewhat similar in spirit to the uniformization step in [17].

5.1 Chain Bound and Chain Algorithm (CA)

Our first proof sequence adapts Jaikumar Radhakrishnan's proof of Shearer's lemma [20] to lattices, which is based on an arbitrary but fixed order of the random variables. We observe that a variable ordering corresponds to a maximal chain in the Boolean algebra. This idea allows us to generalize Radhakrishnan's proof to general lattices. Fix a chain $\mathbf{C} : \hat{0} = C_0 \prec C_1 \prec \cdots \prec C_k = \hat{1}$ in L. For $X \in L$, we say that X *covers* i if $X \wedge C_i \neq X \wedge C_{i-1}$. Intuitively, X covers i if it contains some variable in C_i that does not appear in C_{i-1}. Fix a query represented by $(\mathbf{L}, \mathbf{R})$, where $\mathbf{R} = (R_1, \ldots, R_m)$.

Definition 5.1. The *chain hypergraph* associated with a chain $\mathbf{C}$ is $H_{\mathbf{C}} = ([k], \{e_j \mid j \in [m]\})$, where the hyperedge e_j contains all nodes i such that R_j covers i.

In a Boolean algebra all maximal chains have the same hypergraph, which is the same as the query hypergraph, and the co-atomic hypergraph. But in a general lattice all these hypergraphs may be different, and some output inequalities (7) can be derived only from non-maximal chains. For that reason, we relax the maximality requirement on the chain, as follows. We say that the chain $\mathbf{C}$ is *good* for R_j if:

$$\text{for all } i \in [k]: \quad i \in e_j \Rightarrow C_{i-1} \vee (R_j \wedge C_i) = C_i. \tag{11}$$

The key property of "goodness" is that submodularity applies in the following way:

Proposition 5.2. *If* $\mathbf{C}$ *is a maximal chain, then it is good for any* R_j*. Furthermore, if* $\mathbf{C}$ *is good for* R_j*, then*

$$h(R_j \wedge C_i) - h(R_j \wedge C_{i-1}) \geqslant h(C_i) - h(C_{i-1}), \quad \forall i \in e_j \tag{12}$$

for every $\mathbf{L}$*-submodular function* h*.*

Proof. Notice that $Z \stackrel{\text{def}}{=} C_{i-1} \vee (R_j \wedge C_i)$ is always in the interval $[C_{i-1}, C_i]$. In any chain, $i \notin e_j$ iff $R_j \wedge C_i = R_j \wedge C_{i-1}$ iff $Z = C_{i-1}$. Therefore, any maximal chain is good for R_j. To show (12), replace C_i with $(R_j \wedge C_i) \vee C_{i-1}$ and the inequality becomes the (lattice) submodularity constraint for the elements $(R_j \wedge C_i)$ and C_{i-1}. □

We say a chain $\mathbf{C}$ is *good* for $\mathbf{R}$ if it is good for all $R_j \in \mathbf{R}$. Radhakrishnan's proof is adapted to a lattice as follows.

Theorem 5.3 (The Chain Bound)**.** *Let* $\mathbf{C}$ *be any chain that is good for* $\mathbf{R}$*. If* $(w_j)_{j=1}^m$ *is any fractional edge cover of the chain hypergraph* $H_{\mathbf{C}}$*, then inequality (7) holds for any polymatroid* h

Proof. By writing $h(R_j)$ as a telescoping sum, we obtain

$$\begin{aligned}
\sum_{j=1}^m w_j h(R_j) &= \sum_{j=1}^m w_j \cdot \left(\sum_{i \in e_j} (h(R_j \wedge C_i) - h(R_j \wedge C_{i-1})) \right) \\
\text{(from Eq.(12))} &\geqslant \sum_{j=1}^m w_j \cdot \left(\sum_{i \in e_j} (h(C_i) - h(C_{i-1})) \right) \\
&= \sum_{i=1}^k \left(\sum_{j : i \in e_j} w_j \right) \cdot (h(C_i) - h(C_{i-1})) \\
&\geqslant \sum_{i=1}^k (h(C_i) - h(C_{i-1})) \\
&= h(\hat{1}),
\end{aligned}$$

where the last inequality holds because $(w_j)_{j=1}^m$ is a fractional edge cover of $H_{\mathbf{C}}$ and h is monotone. □

Remark 5.4. Note that, if $\mathbf{C}$ is only good for a subset $\mathbf{R}'$ of $\mathbf{R}$, then we can apply the bound to $\mathbf{R}'$ with its corresponding chain hypergraph.

Example 5.5. Consider our running example from Figure 1; assume $|R| = |S| = |T| = N$. Consider the chain $\hat{0} \prec y \prec yz \prec xyzu = \hat{1}$, whose hypergraph has three vertices and hyperedges $e_R = \{y, xyzu\}$, $e_S = \{y, yz\}$, $e_T = \{yz, xyzu\}$ (isomorphic to the co-atomic hypergraph in Fig.2). Thus, the chain bound on this chain is $N^{3/2}$, which is tight. (Consider the input $R = S = T = [\sqrt{N}] \times [\sqrt{N}]$, where the fd $xz \to u$ is defined by the UDF $f(x, z) = x$ and $yu \to x$ by $g(y, u) = u$.)

If there were no FD's, then the chain bound is exactly Shearer's lemma [8] (or, equivalently, AGM bound).

Corollary 5.6 (AGM bound and Shearer's lemma)**.** *Consider a join query on* n *variables with no fd's. The chain bound on the chain* $\mathbf{C}$*:* $C_0 = \varnothing \prec C_1 = [1] \prec C_2 = [2] \prec \cdots \prec C_n = [n]$ *is exactly Shearer's lemma.*

The Chain Algorithm. In the proof above of the chain bound, the main idea is to take mixtures of conditional entropies, climbing up the chain. This strategy corresponds combinatorially to conditional search. From this proof, we derive Algorithm 1 that computes a query Q in time bounded by *any* fractional edge cover of the chain hypergraph. The algorithm assumes a fixed, good chain $\mathbf{C}$, where every node i is

Algorithm 1 The Chain Algorithm.

Input: A query $(\mathbf{L}, \mathbf{R})$, over $\mathbf{R} = \{R_1, \ldots, R_m\}$
Input: A good chain $\mathbf{C} : \hat{0} = C_0 \prec C_1 \prec \cdots \prec C_k = \hat{1}$
Input: Chain hypergraph $H_{\mathbf{C}} = ([k], \{e_1, \ldots, e_m\})$
1: Expand Q to Q^+ ▷ Thus $\text{vars}(R_j) = R_j$, $j \in [m]$
2: $Q_0 = \{()\}$
3: **For** $i = 1$ **to** k **do**
4: $\quad Q_i = \varnothing$
5: $\quad$ **For** each $t \in Q_{i-1}$ **do**
6: $\quad\quad T = \bigcap_{j : i \in e_j} (t \bowtie \Pi_{R_j \wedge C_i}(R_j))^+$ ▷ Takes time $\tilde{O}\left(\min_{j:i\in e_j} |t \bowtie \Pi_{R_j \wedge C_i}(R_j)|\right)$. See text.
7: $\quad\quad Q_i = Q_i \cup T$
8: **Return** Q_k

covered[7]. It preprocesses input relations by indexing them in an attribute order consistent with the chain. Then, it starts by expanding the query, as explained at the end of Sec. 2. Thus far it takes $\tilde{O}(N)$-time. Next, the algorithm computes inductively

$$Q_i \stackrel{\text{def}}{=} \left(\bowtie_{j : R_j \wedge C_i \neq \hat{0}} \Pi_{R_j \wedge C_i}(R_j) \right)^+,$$

for $i = 0, 1, \ldots, k$, where Q_i is an intermediate relation with attributes C_i. Initially Q_0 consists of just the empty tuple. Evidently, when $i = k$, Q_k is the output Q.

To compute Q_i, let R_j be some relation such that $i \in e_j$, in other words R_j has some new variable that occurs in C_i but not in C_{i-1}. Note that there exists at least one such R_j, because i is covered. Define

$$T_{ij} = Q_{i-1} \bowtie \Pi_{C_i \wedge R_j}(R_j),$$

whose attributes are $X_{ij} \stackrel{\text{def}}{=} C_{i-1} \cup (C_i \wedge R_j)$; and, by Eq.(11), its closure is $X_{ij}^+ = C_{i-1} \vee (C_i \wedge R_j) = C_i$. Consider T_{ij}'s expansion T_{ij}^+ (see Sec. 2): it has the same size as T_{ij} and has attributes $X_{ij}^+ = C_i$. Therefore,

$$Q_i = \bigcap_{j : i \in e_j} T_{ij}^+ = \bigcap_{j : i \in e_j} \left(Q_{i-1} \bowtie \Pi_{C_i \wedge R_j}(R_j) \right)^+.$$

However, we do not want to compute all the T_{ij}^+ and then compute the intersection to obtain Q_i, because this naïve strategy will push the runtime over the budget. In order to stay within the time budget, the algorithm computes this intersection differently: it iterates over all tuples $t \in Q_{i-1}$, and for each tuple computes the intersection T in line 6 in time $\tilde{O}(\min_{j:i\in e_j}(|t \bowtie \Pi_{R_j \wedge C_i}(R_j)|)$. This can be accomplished by first computing $j_* = \text{argmin}_{j:i\in e_j}(|t \bowtie \Pi_{R_j \wedge C_i}(R_j)|)$ and tentatively setting $T = (t \bowtie \Pi_{R_{j_*} \wedge C_i}(R_{j_*}))^+$. Then, the algorithm removes from T any tuple t' that is not in the intersection defined in line 6. A tuple $t' \in T$ is not removed from T only if the following holds: for **every** $j \neq j_*$ s.t. R_j covers i, we have $\Pi_{C_i \wedge R_j}(t') \in \Pi_{C_i \wedge R_j}(R_j)$ **and** $(t \bowtie \Pi_{C_i \wedge R_j}(t'))^+ = t'$. Note the crucial fact that the relation R_{j_*} that is used to iterate over may depend on the tuple $t \in Q_{i-1}$. Due to the pre-processing step where every input relation is indexed with an attribute order consistent with the chain, j_* can easily be computed in logarithmic time in data complexity.

Theorem 5.7. *Assume the chain* $\mathbf{C}$ *is good for* $\mathbf{R}$*, and every node* i *is covered (i.e. no isolated vertices). Then, for any fractional edge cover of the chain hypergraph,* $(w_j)_{j=1}^m$*, the time taken by the Chain Algorithm to compute* Q *is* $\tilde{O}(N + \prod_{j=1}^m N_j^{w_j})$*, where* $\tilde{O}$ *hides a logarithmic factor needed for index lookup or binary search, and a small polynomial factor in query complexity.*

[7] If i is not covered, in other words $H_{\mathbf{C}}$ has an isolated vertex, then $\rho^*(H_{\mathbf{C}}) = \infty$, and the algorithm will not work.

Proof. We define some notation. For any tuple t and each relation R_j such that $R_j \wedge C_i \neq \hat{0}$, denote $n_{ijt} = |t \bowtie \Pi_{R_j \wedge C_i}(R_j)|$. Note that if $t = ()$ is the empty tuple, then $n_{ijt} = n_{ij()} = |\Pi_{R_j \wedge C_i}(R_j)|$.

Fix an iteration i. Denote by Time_i the time taken by iteration i (to compute Q_i). We will show that, for any $i \in [k]$,

$$\text{Time}_i = \tilde{O}\left(\prod_{j:R_j \wedge C_i \neq \hat{0}} n_{ij()}^{w_j}\right). \tag{13}$$

Since $n_{ij()} \leqslant |R_j|$, the sum over all Time_i will be at most $\tilde{O}(\prod_{j=1}^m |R_j|^{w_j})$, as desired. To bound Time_i, note that the number of steps taken to compute Q_i is

$$\sum_{t \in Q_{i-1}} \min_{j:i \in e_j} n_{ijt} \leqslant \sum_{t \in Q_{i-1}} \prod_{j:i \in e_j} n_{ijt}^{w_j} = \sum_{t \in Q_{i-1}} \prod_{j:R_j \wedge C_i \neq \hat{0}} n_{ijt}^{w_j}.$$

The equality follows from the fact that, if $t \in Q_{i-1}$ and $R_j \wedge C_i \neq \hat{0}$ but $i \notin e_j$, then $n_{ijt} = 1$. Hence, to show (13), it is sufficient to show the following:

$$\sum_{t \in Q_{i-1}} \prod_{j:R_j \wedge C_i \neq \hat{0}} n_{ijt}^{w_j} \leqslant \prod_{j:R_j \wedge C_i \neq \hat{0}} |n_{ij()}|^{w_j}. \tag{14}$$

For $0 \leqslant \ell \leqslant i-1$, define

$$\begin{aligned} T_\ell &\stackrel{\text{def}}{=} \Pi_{C_\ell}(Q_{i-1}), \\ g(\ell) &\stackrel{\text{def}}{=} \sum_{t \in T_\ell} \prod_{j:R_j \wedge C_i \neq \hat{0}} n_{ijt}^{w_j}. \end{aligned}$$

By convention, T_0 has a single empty tuple $t = ()$. Then, (14) is equivalent to $g(i-1) \leqslant g(0)$. Thus, to show (14) it is sufficient to show that $g(\ell)$ is non-increasing in ℓ:

$$\begin{aligned} g(\ell) &= \sum_{t \in T_\ell} \prod_{j:R_j \wedge C_i \neq \hat{0}} n_{ijt}^{w_j} \\ &= \sum_{u \in T_{\ell-1}} \sum_{\substack{v: \\ (u,v) \in T_\ell}} \prod_{j:R_j \wedge C_i \neq \hat{0}} n_{ij(u,v)}^{w_j} \\ &= \sum_{u \in T_{\ell-1}} \prod_{\substack{j:R_j \wedge C_i \neq \hat{0} \\ i-1 \notin e_j}} n_{iju}^{w_j} \sum_{\substack{v: \\ (u,v) \in T_\ell}} \prod_{\substack{j:R_j \wedge C_i \neq \hat{0} \\ i-1 \in e_j}} n_{ij(u,v)}^{w_j} \\ &\leqslant \sum_{u \in T_{\ell-1}} \prod_{\substack{j:R_j \wedge C_i \neq \hat{0} \\ i-1 \notin e_j}} n_{iju}^{w_j} \prod_{\substack{j:R_j \wedge C_i \neq \hat{0} \\ i-1 \in e_j}} \left(\sum_{\substack{v: \\ (u,v) \in T_\ell}} n_{ij(u,v)}\right)^{w_j} \\ &\leqslant \sum_{u \in T_{\ell-1}} \prod_{\substack{j:R_j \wedge C_i \neq \hat{0} \\ i-1 \notin e_j}} n_{iju}^{w_j} \prod_{\substack{j:R_j \wedge C_i \neq \hat{0} \\ i-1 \in e_j}} n_{iju}^{w_j} \\ &= \sum_{u \in T_{\ell-1}} \prod_{j:R_j \wedge C_i \neq \hat{0}} n_{iju}^{w_j} \\ &= g(\ell-1). \end{aligned}$$

The first inequality is generalized Hölder inequality, which applies because $(w_j)_{j=1}^m$ fractionally cover vertex $i-1$ of the chain hypergraph $H_{\mathbf{C}}$. The second inequality holds because, for every $u \in T_{\ell-1}$, $u \bowtie \Pi_{R_j \cap C_i}(R_j) \supseteq T_\ell$. ∎

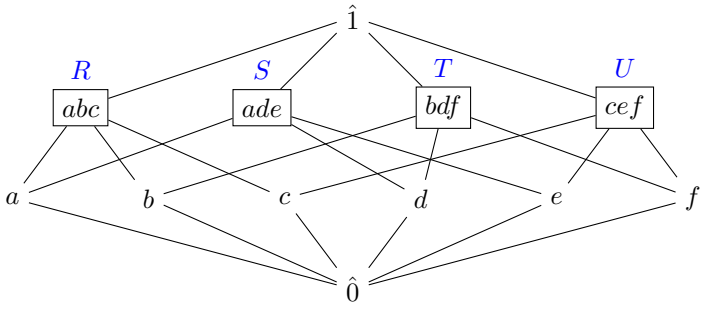

Figure 4: A query where the chain bound is not optimal

We present some examples.

Example 5.8. Continue with Example 5.5. The variable order corresponding to the chain $\hat{0} \prec y \prec yz \prec \hat{1}$ is $y, z, (xu)$, where x, u can be arranged in any order. The Chain Algorithm computes three intermediate relations:

$$\begin{aligned} Q_1(y) &= \Pi_y(R(xy)) \cap \Pi_y(S(yz)) \\ Q_2(yz) &= Q_1(y) \bowtie S(yz) \\ Q_3(yzxu) &= (Q_2(yz) \bowtie R(xy))^+ \cap (Q_2(yz) \bowtie T(zu))^+ \end{aligned}$$

The first two steps are straightforward. In the third, the join $Q_2(yz) \bowtie R(xy)$ results in a relation with attributes xyz, which needs to be expanded with u (e.g. by computing the user-defined function $u = f(x, z)$), and similarly for the second join. However, the algorithm does not compute the joins first then intersect, instead it iterates over tuples $t \in Q_2(yz)$ and computes an intersection on a per-tuple basis using the less expensive option. In particular, for each $t = (y, z) \in Q_2$, it compares $|t \bowtie R(x, y)|$ and $|t \bowtie T(z, u)|$ (which can be done in logarithmic time given that R was indexed with attribute order (y, x) and T with order (z, u)). Suppose $|t \bowtie R(x, y)| \leqslant |t \bowtie T(z, u)|$, then for each $(y, z, x) \in t \bowtie R(x, y)$ the algorithm uses the $xz \to u$ FD to obtain the tuple (y, z, x, u). The next task is to verify that this tuple is indeed in the intersection as defined in line 6. This is done with *two* verifications: we make sure that $(z, u) \in T$, **and** that $yu \to x$ is indeed satisfied.[8]

On this chain the algorithm runs in optimal time $O(N^{3/2})$. We note that all previously proposed known worst-case optimal join algorithms for queries without FD's [18,19,23] require $\Omega(N^2)$ to compute the previous query on this instance: $R = S = T = \{(1, i) \mid i \in [N/2]\} \cup \{(i, 1) \mid i \in [N/2]\}$. For example, LFTJ with variable order y, z, x, u computes queries $Q_1(y), Q_2(yz), Q_3(xyz), Q_4(xyzu)$, where $|Q_3| = N^2$. Note, however, that not every maximal chain gives an optimal bound: for example the chain $\hat{0} \prec x \prec xu \prec xyu \prec xyzu = \hat{1}$ has hyperedges: $e_R = \{x, xyu\}$, $e_S = \{xyu, xyzu\}$, $e_T = \{xu, xyzu\}$ (isomorphic to the atomic hypergraph in Fig.2), and the optimal fractional edge covering number $\rho^* = 2$, hence the chain bound is $|Q^D| \leqslant N^2$, which is sub-optimal.

Choosing a good chain. When there are FDs, it is not clear if there even exists a chain with a finite chain bound. We show in this section how to select a chain so that the chain hypergraph does not have an isolated vertex. This means the fractional edge cover number of the chain hypergraph is finite. The first method to select the chain corresponds to a generalization of Shearer's lemma. The second method interestingly corresponds to a *dual* version of Shearer's lemma.

Corollary 5.9 (Shearer's Lemma for FDs)**.** *Consider a query* $(\mathbf{L}, \mathbf{R})$*. Let* $\mathbf{J}$ *be the set of join-irreducibles that are below the order ideal generated by elements in* $\mathbf{R}$*. There exists a set* $\{X_1, \ldots, X_n\}$ *of members of* $\mathbf{J}$

[8]This is a subtle step in the algorithm that is easy to miss at the first read.

satisfying the following. Let $\mathbf{C}$ *be the chain* $C_0 = \hat{0} \prec C_1 \prec \cdots \prec C_n = \hat{1}$*, where* $C_i = \bigvee_{j=1}^{i} X_j$*. Then the chain* $\mathbf{C}$ *is good and the hypergraph* $H_{\mathbf{C}}$ *has no isolated vertex.*

Proof. Note that $\bigvee_{Z \in \mathbf{J}} Z = \hat{1}$ because $\bigvee_{j=1}^{m} R_j = \hat{1}$. We construct the sequence $X_1, \ldots, X_n$ inductively. For $i \in [n]$, suppose we have constructed the sequence $X_1, \ldots, X_{i-1}$. Let $Y = \bigvee_{j=1}^{i-1} X_j$. (If $i = 1$ then $Y = \hat{0}$.) Let X_i be an element in $\mathbf{J} - \{X_1, \ldots, X_{i-1}\}$ such that the $Y \prec Y \vee X_i$ and $Y \vee X_i$ is *minimal* among all such X_i. We stop at X_n when the join is $\hat{1}$.

For $i \in [n]$, let $C_i = \bigvee_{j=1}^{i} X_j$. We first show that the chain $\mathbf{C} = (C_i)_{i=1}^{n}$ is a good chain for $\mathbf{R}$. Note that all elements C_i are distinct. Consider any $R_j \in \mathbf{R}$, and suppose $R_j \wedge C_{i-1} \neq R_j \wedge C_i$ but $(R_j \wedge C_i) \vee C_{i-1} \prec C_i$. Then, $S = (R_j \wedge C_i) \vee C_{i-1}$ is strictly in between C_{i-1} and C_i. Because, if $S = C_{i-1}$ then $R_j \wedge C_{i-1} = R_j \wedge C_i$. Let $Y = \bigvee_{j=1}^{i-1} X_j$. We show that there is an element $X \in \mathbf{J} - \{X_1, \ldots, X_{i-1}\}$ such that $C_{i-1} \prec Y \vee X \prec C_i$. This will violate the choice of X_i and the proof of the claim would be complete. Note that S is the join of all join-irreducibles below C_{i-1} and below $R_j \wedge C_i$. The join-irreducibles below $R_j \wedge C_i$ are below R_j, and thus they are in the set $\mathbf{J}$. Pick X to be a join-irreducible below $R_j \wedge C_i$ but not below C_{i-1} and we are done.

To see that every vertex $i \in [n]$ is covered in the chain-cover hypergraph, note that if $X_i \preceq R_j$ then $R_j \wedge C_i \neq R_j \wedge C_{i-1}$. ☐

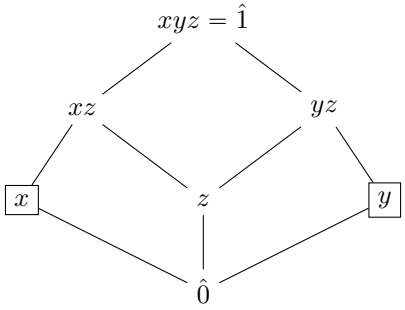

Figure 5: Lattice for Q :- $R(x), S(y), xy \to z$

Example 5.10. Recall an example query with UDFs: Q :- $R(x), S(y), z = f(x, y)$ where f is a UDF. The lattice is shown in Fig. 5. If we selected any maximal (and thus good) chain, such as $\hat{0} \prec z \prec xz \prec xyz$, or $\hat{0} \prec x \prec xz \prec xyz$, then z or xz would be an isolated vertex in the chain hypergraph. Corollary 5.9 tells us to construct a chain by joining the join irreducibles below R and S, which are x, y. Hence, we would select a chain such as $\hat{0} \prec x \prec xyz$ which has no isolated vertices. The algorithm runs in time $O(N^2)$, which is worst-case optimal. This chain is not maximal.

We can flip the above proof, working from the meet-irreducibles instead of the join-irreducibles, to obtain the dual version of Shearer's lemma.

Corollary 5.11 (Dual Shearer's Lemma for FDs)**.** *Let* $(\mathbf{L}, \mathbf{R})$ *represent an input query. There exists a sequence* $X_1, \ldots, X_n$ *of meet-irreducibles of* $\mathbf{L}$ *satisfying the following. For* $i = 0, 1, \ldots n$*, define* $C_i = \bigwedge_{j=1}^{n-i} X_j$*. The chain* $\mathbf{C} = (C_i)_{i=0}^{n}$ *is good for* $\mathbf{R}$ *and the chain hypergraph* $H_{\mathbf{C}}$ *has no isolated vertex.*

A condition for the chain bound to be tight. Now that we know how to select a chain so that the bound is finite, the next question is whether the chain bound is tight for some class of queries.

Example 5.12. The chain bound is optimal on some non-normal lattices. Consider the query

$$R(x), S(y), T(z), xy \to z, xz \to y, yz \to x,$$

whose lattice is M_3 in Fig. 3 (a non-normal lattice), with $|R| = |S| = |T| = N$. The chain $\hat{0} \prec x \prec xyz = \hat{1}$ gives a the tight upper bound N^2, because its chain hypergraph is $e_R = \{x\}, e_S = \{\hat{1}\}, e_T = \{\hat{1}\}$, has optimal edge covers $(w_x, w_y, w_z) = (1, 1, 0)$ and $= (1, 0, 1)$.

Before presenting a sufficient condition for when the chain bound is tight, we need a simple lemma.

Lemma 5.13. *Let $(\mathbf{L}, \mathbf{R})$ represent a query. Let $\mathbf{C} : \hat{0} = C_0 \prec C_1 \prec \cdots \prec C_k = \hat{1}$ be a chain that is good for every $X \in L$. For every $S \in L$, define $e(S) = \{i \in [k] \mid S \wedge C_i \neq S \wedge C_{i-1}\}$. Then, $X \preceq Y$ implies $e(X) \subseteq e(Y)$.*

Proof. Consider any $i \notin e(Y)$; then, $X \wedge C_i \preceq Y \wedge C_i = Y \wedge C_{i-1} \preceq C_{i-1}$. Thus, $X \wedge C_i = X \wedge C_{i-1}$ which means $i \notin e(X)$. □

Theorem 5.14. *Let $\hat{0} = C_0 \prec C_1 \cdots \prec C_k = \hat{1}$ be a chain that is good for every $X \in L$. Also, suppose for every $X, Y \in L$, the following holds*

$$e(X \vee Y) \subseteq e(X) \cup e(Y). \tag{15}$$

Then the chain bound is tight on this lattice.

Proof. Let h^* be an optimal (polymatroid) solution to LLP on L. We define a new $\mathbf{L}$-function u as follows.

$$\begin{aligned} u(\hat{0}) &= 0 \\ u(X) &= \sum_{i \in e(X)} (h^*(C_i) - h^*(C_{i-1})),\ X \in L. \end{aligned}$$

Claim 1. u is also an optimal solution to LLP on $\mathbf{L}$.

We first show that u is a polymatroid. Non-negativity of u follows from monotonicity of h^*. Monotonicity of u follows from Lemma 5.13. Submodularity follows from the assumption that $e(X \vee Y) \subseteq e(X) \cup e(Y)$. In fact, from Lemma 5.13 $e(X) \subseteq e(X \vee Y)$ and $e(Y) \subseteq e(X \vee Y)$, which means $e(X \vee Y) = e(X) \cup e(Y)$. And thus u is modular.

Next, we show that $u(X) \leqslant h^*(X)$ for all $X \in L$. In particular, u is a feasible solution to LLP on L. This is proved by induction on $|e(X)|$. The base case when $|e(X)| = 0$ is trivial. For the inductive step, let j be the maximum number in $e(X)$. Then,

$$\begin{aligned} u(X) &= \sum_{i \in e(X)} (h^*(C_i) - h^*(C_{i-1})) \\ &= \sum_{i \in e(X \wedge C_{j-1})} (h^*(C_i) - h^*(C_{i-1})) + h^*(C_j) - h^*(C_{j-1}) \\ &= u(X \wedge C_{j-1}) + h^*(C_j) - h^*(C_{j-1}) \\ \text{(induction hypothesis)} &\leqslant h^*(X \wedge C_{j-1}) + h^*(C_j) - h^*(C_{j-1}) \\ \text{(submodularity of } h^*) &\leqslant h^*(X) \end{aligned}$$

Since $u(\hat{1}) = h^*(\hat{1})$, u is an optimal solution to LLP on L as claimed.

Claim 2. u is materializable.

Let g be a polymatroid on the Boolean algebra B_k, where

$$\begin{aligned} g(i) &= h^*(C_i) - h^*(C_{i-1}) \\ g(X) &= \sum_{i \in X} g(i). \end{aligned}$$

Then g is a modular polymatroid and it can be materialized with a product instance: $D = \prod_{i=1}^{k} [2^{g(i)}]$. We exhibit an embedding from u to f. The map $f : L \to B_k$ defined by $f(X) = e(X)$ is an embedding because $X \vee Y \to e(X \vee Y) = e(X) \cup e(Y)$. One can verify that $u = g \circ f$. Hence, u is materializable from Proposition 3.6. □

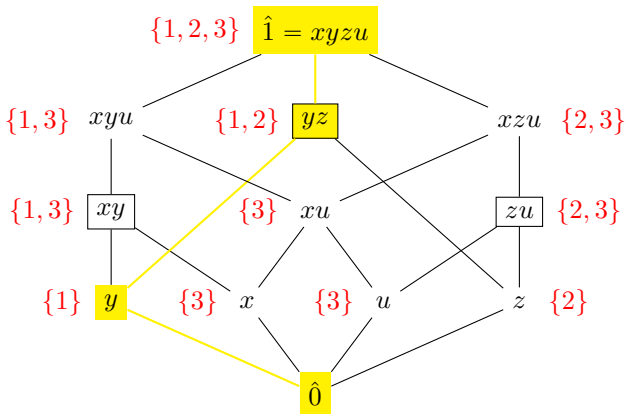

Figure 6: Condition (15) holds for this lattice

Corollary 5.15. *The chain bound is tight on distributive lattices.*

Proof. Consider any maximal chain on the distributive lattice L. We only have to verify that

$$e(X \vee Y) \subseteq e(X) \cup e(Y).$$

We prove this by showing that if $i \notin e(X) \cup e(Y)$ then $i \notin e(X \vee Y)$. Suppose

$$\begin{aligned} X \wedge C_i &= X \wedge C_{i-1} \\ Y \wedge C_i &= Y \wedge C_{i-1}. \end{aligned}$$

Then,

$$\begin{aligned} (X \vee Y) \wedge C_i &= (X \wedge C_i) \vee (Y \wedge C_i) \\ &= (X \wedge C_{i-1}) \vee (Y \wedge C_{i-1}) \\ &= (X \vee Y) \wedge C_{i-1}. \end{aligned}$$

□

Example 5.16 (Tightness on non-distributive lattice)**.** The chain bound is tight on the lattice (and the chain) shown in Figure 6. The red sets are the sets $e(X)$, $X \in L$. In particular, the characterization condition (15) goes beyond distributive lattices.

From Proposition 3.2 we obtain the following which subsumes results from [11] in the simple FD case.

Corollary 5.17. *If all FDs are simple, then the chain bound is tight and the chain algorithm is worst-case optimal.*

Example 5.18 (Chain bound is not always tight)**.** For some queries, even with normal lattices, no chains give a tight upper bound. This is illustrated by Q in Fig. 4. Consider the chain $\hat{0} \prec a \prec abc \prec \hat{1}$. Denoting $\{1,2,3\}$ the vertices of the chain hypergraph, its edges are $e_R = \{1,2\}$, $e_S = \{1,3\}$, $e_T = e_U = \{2,3\}$. The optimal fractional edge cover is $(w_R, w_S, w_T, w_U) = (1/2, 1/2, 1/2, 0)$, $\rho^* = 3/2$, therefore the chain bound is $|Q^D| \leqslant N^{3/2}$ and the Chain Algorithm will run in this time. All other maximal chains give the same bound (non-maximal chains are not good). However, we show in the next section that $|Q^D| \leqslant N^{4/3}$, which means that no chain bound is tight.

Discussion We have shown that the chain bound implies Shearer's lemma and AGM bound from the Boolean algebra; more generally it is tight on all distributive lattices. The lattice corresponding to simple FDs is distributive, thus the chain algorithm is worst-case optimal for simple FDs. We also presented examples and proved results on how to choose a good chain so that the chain hypergraph does not have an isolated vertex.

For brevity, we have described the Chain Algorithm using breadth-first search (or bottom-up). It can also be adapted to a depth-first implementation (or top-down), as used for example in LogicBlox' LFTJ [23], which does not materialize intermediate relations.

5.2 Sub-Modularity Bound and Sub-Modularity Algorithm (SMA)

For some lattices, like Example 5.18, the Chain Algorithm is sub-optimal no matter what chain we choose. We describe here a second proof technique for Shearer's lemma and its adaptation to lattices, and derive a new algorithm. We assume w.l.o.g. that all coefficients w_j in (7) are rational numbers, $w_j = q_j/d$, and rewrite the set $\{R_1, \ldots, R_m\}$ of lattice elements as a multiset $\mathcal{B} = \left\{B_1, B_2, \ldots, B_{\sum_j q_j}\right\} \subseteq L$ where each lattice element R_j is copied q_j times. Then, inequality (7) becomes:

$$\sum_i h(B_i) \geqslant d \cdot h(\hat{1}) \tag{16}$$

The Sub-modularity Proof Sequence. Balister and Bollobás [4] give a simple proof of (16) that uses only the sub-modularity inequality, which we adapt here for arbitrary lattices. We refer to this proof strategy a *sub-modularity proof* or SM-proof. An SM-proof starts with a mutiset $\mathcal{B} = \{B_1, B_2, \ldots\}$, and applies repeatedly *sub-modularity steps* (SM-steps)[9]. An SM-step, $(X, Y) \rightarrow (X \wedge Y, X \vee Y)$, consists of removing two incomparable elements X, Y from $\mathcal{B}$, and replacing them with $X \wedge Y, X \vee Y$. A *sub-modularity proof sequence*, or simply an SM-proof, repeatedly applies SM-steps to a multiset $\mathcal{B}$ until all elements in $\mathcal{B}$ are comparable; at that point $\mathcal{B}$ is a chain $\hat{1} \succeq C_1 \succeq C_2 \succeq \ldots$, and we denote $d, d_1, d_2, \ldots$ the multiplicities of the elements in $\mathcal{B}$. The *sub-modularity bound* corresponding to this proof sequence is:

$$\sum_i h(B_i) \geqslant d \cdot h(\hat{1}) + d_1 \cdot h(C_1) + \ldots \geqslant d \cdot h(\hat{1})$$

Example 5.19. The following is a valid instance of inequality (7) on the Boolean algebra lattice $2^{\{a,b,c,d\}}$:

$$h(abc) + \frac{1}{3}h(abd) + \frac{1}{3}h(acd) + \frac{1}{3}h(bcd) \geqslant h(\hat{1}). \tag{17}$$

This corresponds to $d = 3$, $\mathcal{B} = \{abc, abc, abc, abd, acd, bcd\}$, and the rewritten inequality

$$h(abc) + h(abc) + h(abc) + h(abd) + h(acd) + h(bcd) \geqslant 3 \cdot h(\hat{1}).$$

We can prove this inequality by applying the following SM-steps:

$$\begin{aligned}
h(abc) + h(abd) &\geqslant h(\hat{1}) + h(ab) \\
h(abc) + h(acd) &\geqslant h(\hat{1}) + h(ac) \\
h(abc) + h(bcd) &\geqslant h(\hat{1}) + h(bc) \\
h(ab) + h(ac) &\geqslant h(abc) + h(a) \\
h(a) + h(bc) &\geqslant h(abc) + h(\hat{0}).
\end{aligned}$$

Altogether the SM-steps proved that

$$h(abc) + h(abc) + h(abc) + h(abd) + h(acd) + h(bcd) \geqslant 3 \cdot h(\hat{1}) + 2h(abc) + h(\hat{0}),$$

[9] Called "elementary compression" in [4].

which implies (17). This example also shows that there is extra information we could not make use of: the $2h(abc)$ term that is left "dangling". As we shall see later, in some cases this is a manifestation of the limitation of SM-proofs.

Example 5.20. Continuing Example 5.18 where the Chain bound is not tight, the SM-proof is very simple:

$$\begin{aligned} h(abc) + h(ade) &\geqslant h(\hat{1}) + h(a) \\ h(bdf) + h(cef) &\geqslant h(\hat{1}) + h(f) \\ h(a) + h(f) &\geqslant h(\hat{1}) + h(\hat{0}), \end{aligned}$$

resulting in the SM-bound:

$$h(\hat{1}) \leqslant \frac{1}{3}h(abc) + \frac{1}{3}h(ade) + \frac{1}{3}h(bdf) + \frac{1}{3}h(cef).$$

In particular, when all input relations have size N, the output size bound is $N^{4/3}$. This coincides with the co-atomic hypergraph cover, hence it is tight.

Obviously, any SM-bound is a correct inequality of the form (16), but the converse does not always hold, as we shall see. The converse holds, however, in distributive lattices. We next show that given any fractional edge cover $(w_j)_{j=1}^m$ of the co-atomic hypergraph of a distributive lattice, inequality (7) is provable through an SM-proof sequence and, moreover, one can apply the SM-steps in any order. It follows that every distributive lattice is normal!

For each element $X \in L$, define $e_X = \{Z \mid Z \in V_{\text{co}}, X \not\leq Z\}$; notice that, if X is an input, $X = R_j$, then e_{R_j} is a hyperedge of H_{co}. A d-cover of the co-atoms is a multiset $\mathcal{B}$ s.t. for each vertex $Z \in V$ there are at least d elements $X \in B$ s.t. $Z \in e_X$. Then:

Lemma 5.21. *Suppose L is a distributive lattice, and B is a d-cover of the co-atoms. If B' is obtained from B by applying one SM-step $(X, Y) \to (X \wedge Y, X \vee Y)$, then B' is also a d-cover of the co-atoms.*

Proof. We prove this using a similar argument as in [4]. By Lemma 4.1, for any co-atom Z, $c(Z, \{X, Y\}) = c(Z, \{X \wedge Y, X \vee Y\})$, hence the number of elements in B that cover Z is the same as the number of elements in B' that cover Z. □

Corollary 5.22. *Given any fractional edge cover $(w_j)_{j=1}^m$ of the co-atomic hypergraph, inequality (7) is provable through an SM-proof sequence and, moreover, one can apply the SM-steps in any order.*

Proof. Progress is ensured by the fact that $\sum_i |e_{B_i}|^2$ strictly increases after each sub-modularity step:

$$|e_X|^2 + |e_Y|^2 < |e_X \cap e_Y|^2 + |e_X \cup e_Y|^2 = |e_{X \vee Y}|^2 + |e_{X \wedge Y}|^2.$$

Finally, when the process ends, the multiset $\mathcal{B}$ is a chain and $h(\hat{1})$ must occur d times, because any co-atom Z that is $\succeq C_1$ (the next largest element in the chain) is covered only by $\hat{1}$. □

Corollary 5.23. *Every distributive lattice is normal.*

The Submodularity Algorithm. The SM-Algorithm (Algorithm 2) starts by using the cardinalities $(N_j)_{j=1}^m$ of the input relations to obtain an optimal solution h^* of the LLP (Eq.(5)), and an optimal dual solution s^*, w^*: the coefficients $(w_j^*)_{j=1}^m$ form a valid inequality (7) (see the discussion at the beginning of Sec 5), which is tight for h^*; as before, write $s_{X,Y}^*, w_j^*$ as rational numbers, $w_j^* = q_j/d, s_{X,Y}^* = p_{X,Y}/d$.[10] The SM-algorithm requires as an input an SM-proof sequence of this inequality, then computes the query Q^D as follows.

[10] Extreme points of the dual polytope are data-independent!

The algorithm performs the SM-steps in the proof sequence, maintaining the multiset $\mathcal{B} \subseteq L$.[11] It also maintains a cache of temporary relations, in one-to-one correspondence with $\mathcal{B}$: for each $B \in \mathcal{B}$, there is one temporary relation $T(B)$, with set of attributes B; if B occurs multiple times in $\mathcal{B}$, then there are multiple temporary relations $T(B)$. Initially the temporary tables are the input relations, s.t. each relation R_j is copied q_j times. Next, the algorithm applies the SM-steps in the proof sequence, and for each step $(X, Y) \to (X \wedge Y, X \vee Y)$ performs a *sub-modularity join*:

> **Sub-modularity join.** Let $Z = X \wedge Y$ be the set of common variables in $T(X), T(Y)$. The *degree* of some value $v \in \prod_{z \in Z} \mathsf{Domain}(Z)$ is the number of tuples in $T(Y)$ with $Z = v$:
>
> $$\mathsf{degree}_{T(Y)}(v) \stackrel{\text{def}}{=} |\sigma_{Z=v}(T(Y))| \tag{18}$$
>
> The SM-join partitions $\Pi_Z(T(Y))$ into *light hitters* and *heavy hitters*, consisting of values with degree $\leqslant 2^{h^*(Y)-h^*(Z)}$ or $> 2^{h^*(Y)-h^*(Z)}$ respectively; denote them Lite and Heavy. Define $T(X \vee Y)$ the subset of the join $T(X) \bowtie T(Y)$ restricted to light hitters, and define $T(X \wedge Y)$ to be the set of heavy hitters. The SM-join removes $T(X), T(Y)$ from the cache, and adds $T(X \wedge Y), T(X \vee Y)$ to the cache of temporary tables.

After processing the entire proof sequence, the algorithm returns the union of all d temporary tables $T(\hat{1})$, then semi-join reduces them with all input relations.

Algorithm 2 The Sub-modularity Algorithm

Input: A query Q, over relations $R_1, \ldots, R_m$
Input: A SM-proof sequence of $\sum_{B \in \mathcal{B}} h^*(B) \geqslant d \cdot h^*(\hat{1})$
Initialize $|\mathcal{B}|$ temporary tables $T(B)$, $B \in \mathcal{B}$:
  where each $T(B)$ is initially some R_j ▷ See text
For each SM-step $(X, Y) \to (X \wedge Y, X \vee Y)$ **do**
  Let $Z = X \wedge Y$
  $\mathsf{Lite}(Z) \stackrel{\text{def}}{=} \left\{ v \mid \log \mathsf{degree}_{T(Y)}(v) \leqslant h^*(Y) - h^*(Z) \right\}$
  $\mathsf{Heavy}(Z) \stackrel{\text{def}}{=} \left\{ v \mid \log \mathsf{degree}_{T(Y)}(v) > h^*(Y) - h^*(Z) \right\}$
  Add the following tables to the cache
  $T(X \wedge Y) = \Pi_Z(T(X)) \cap \Pi_Z(T(Y)) \cap \mathsf{Heavy}(Z)$
  $T(X \vee Y) = (T(X) \bowtie (T(Y) \ltimes \mathsf{Lite}(Z)))^+$
Return $\bigcup_{B \in \mathcal{B}, B = \hat{1}} T(B)$ semi-join reduced with all inputs

The following invariant is maintained by the algorithm:

Lemma 5.24. *At each step in the algorithm, for any relation $T(B)$ in cache, $\log |T(B)| \leqslant h^*(B)$.*

Proof. We induct on the number of SM-steps. If $T(B)$ was a copy of an input relation R_j that participates in the proof, then $\log |T(B)| = n_j = h^*(B)$, because due to complementary slackness $w_j^* > 0$ implies the primal constraint has to be tight. For every SM-step $(X, Y) \to (X \wedge Y, X \vee Y)$, it must hold that $h^*(Y) - h^*(X \wedge Y) = h^*(X \vee Y) - h^*(X)$ because inequality (7) is an *equality* for h^*: we started from an LLP-dual optimal solution. Hence,

$$\begin{aligned}
|T(X \vee Y)| &\leqslant |T(X)| \cdot 2^{h^*(Y)-h^*(Z)} \\
&\leqslant 2^{h^*(X)} \cdot 2^{h^*(X \vee Y)-h^*(X)} \\
&\leqslant 2^{h^*(X \vee Y)} \\
|T(X \wedge Y)| &\leqslant |T(Y)|/2^{h^*(Y)-h^*(Z)} \leqslant 2^{h^*(Z)}
\end{aligned}$$

□

[11] I.e. every member of the multiset is a member of the set L

Example 5.25. Consider the SM proof sequence from Example 5.20, we explain how the SM algorithm works for this proof sequence. The optimal solution LLP solution is

$$\begin{aligned} h^*(\hat{1}) &= (4/3)\cdot \log N, \\ h^*(X) &= \log N, \quad \text{for } X \in \{abc, ade, bdf, cef\}, \\ h^*(X) &= (2/3)\cdot \log N, \quad \text{for } X \in \{a, b, c, d, e, f\}. \end{aligned}$$

The SM algorithm works as follows.

- It first SM-joins $R(abc)$ with $S(ade)$, producing relations $Q_1(\hat{1})$ and $\mathsf{Heavy}_1(a)$, where $\mathsf{Heavy}_1(a)$ is the set of values a whose S-degree is at least $N^{1-2/3} = N^{1/3}$. Thus, $|\mathsf{Heavy}_1(a)| \leqslant N^{2/3}$. Since the light part has degree at most $N^{1/3}$, $|Q_1(\hat{1})| \leqslant N^{4/3}$.
- Similarly it SM-joins $T(bdf)$ with $U(cef)$ producing relations $Q_2(\hat{1})$ and $\mathsf{Heavy}_2(f)$, and finally it computes the cross product $\mathsf{Heavy}_1(a) \times \mathsf{Heavy}_2(f)$, then expands the result (since $\{a, f\}^+ = \hat{1}$) to obtain a relation $Q_3(\hat{1})$. Since both $|\mathsf{Heavy}_1(a)| \leqslant N^{2/3}$ and $|\mathsf{Heavy}_2(f)| \leqslant N^{2/3}$, their cross-product has size at most $N^{4/3}$ and can be computed within that time budget.
- Finally, it returns $Q_1 \cup Q_2 \cup Q_3$ semi-joined with input relations.

There are two reasons why the SM-algorithm may fail: some branches of heavy or light elements may never join into a $T(\hat{1})$, and the algorithm may attempt to join some light with heavy values from the same relation. We give a sufficient condition for the SMA algorithm to be correct. The criterion consists of iterating over the SM-proof sequence, and assigning a set $\mathsf{Labels}(B)$ of labels to each copy $B \in \mathcal{B}$. (Copies of the same lattice element receive their own label sets.) Initially, every $B \in \mathcal{B}$ receives a single label 1, namely $\mathsf{Labels}(B) = \{1\}$, the same for all $B \in \mathcal{B}$. Consider an SM-step $(X, Y) \to (X \wedge Y, X \vee Y)$. Let $\mathcal{A}(X, Y) = \mathsf{Labels}(X) \cap \mathsf{Labels}(Y)$. Assign $\mathsf{Labels}(X \vee Y) = \mathcal{A}(X, Y)$ and, if $X \wedge Y \neq \hat{0}$, assign to $\mathsf{Labels}(X \wedge Y)$ a *fresh* set of a labels: $\mathsf{Labels}(X \wedge Y) = \{f(j) \mid j \in \mathcal{A}(X, Y)\}$ for some label assignment f. For each $Z \notin \{X, Y\}$, set $\mathsf{Labels}(Z) = \mathsf{Labels}(Z) \cup \{f(j) \mid \mid j \in \mathsf{Labels}(Z) \cap \mathcal{A}(X, Y)\}$. Note the crucial fact that in this description we did not remove X, Y from $\mathcal{B}$. We keep accumulating elements to $\mathcal{B}$, unlike in the proof sequence where each step replaces an old pair with a new pair. Note the important fact that that each copy of lattice element in the multiset $\mathcal{B}$ gets its own label set.

Definition 5.26 (Good SM-proof sequence)**.** Call the SM-proof sequence *good* if $\mathcal{A}(X, Y) \neq \varnothing$ for all SM-steps $(X, Y) \to (X \vee Y, X \wedge Y)$, *and* if in the end every label is present in $\bigcup_{\hat{1} \in B} \mathsf{Labels}(\hat{1})$.

Example 5.27. Continue with Example 5.25, we check that the SM-proof is good. Initially all elements in $\{abc, ade, bdf, cef\}$ have $\mathsf{Labels} = \{1\}$. After the first SM-step: $\mathsf{Labels}(\hat{1}) = \{1\}$ and $\mathsf{Labels}(a) = \{2\}$, where 2 is a fresh label; 2 is also added to Labels of abc, ade, bdf, cef, so they are all equal to $\{1, 2\}$. After the second SM-step there are two copies of $\hat{1}$, where $\bigcup_{\hat{1} \in \mathcal{B}} \mathsf{Labels}(\hat{1}) = \{1, 2\}$ and $\mathsf{Labels}(f) = \{3, 4\}$, two fresh labels, which are added to $\mathsf{Labels}(a)$ so that $\mathsf{Labels}(a) = \{2, 3, 4\}$. After the third SM-step, $\bigcup_{\hat{1} \in \mathcal{B}} \mathsf{Labels}(\hat{1}) = \{1, 2, 3, 4\}$; it follows that the proof sequence is good.

We next show that if there is a good proof sequence, then SMA runs in time that matches the bound.

Theorem 5.28. *If the SM-proof sequence for $\sum_j w_j^* h(R_j) \geqslant h(\hat{1})$ is good, then the SM algorithm correctly computes Q^D, and runs in time $\tilde{O}(N + \prod_j N_j^{w_j^*})$, where $\tilde{O}$ hides a $\log N$ factor, a polynomial in query size and SM-proof length.*

Proof. At any point in time, let $L_i = \{B \in \mathcal{B} \mid i \in \mathsf{Labels}(B)\}$. Let $Q_i \stackrel{\text{def}}{=} \bowtie_{B \in L_i} T(B)$. We show by induction on the number of SM-steps that the algorithm maintains the following invariant: $Q \subseteq \bigcup_i Q_i$. The base case holds because initially there is only one label, and $Q = Q_1$. Consider an SM-step, $(X, Y) \to (X \wedge Y, X \vee Y)$.

Only the subqueries Q_i for which $i \in \mathcal{A}(X, Y)$ contain both X and Y. For each such query, consider two new queries:

$$\begin{aligned} Q_i^{\mathsf{Heavy}} &= T(X \wedge Y) \bowtie\!\!\bowtie_{B \in L_i} T(B) \\ Q_i^{\mathsf{Lite}} &= T(X \vee Y) \bowtie\!\!\bowtie_{B \in L_i} T(B). \end{aligned}$$

Then, obviously $Q_i \subseteq Q_i^{\mathsf{Heavy}} \cup Q_i^{\mathsf{Lite}}$ and thus the invariant is maintained after the SM-step. In the end, each sub-query reaches $\hat{1}$ and the semi-join reduction filters the result. (We need this step because large input relations R_j might have had coefficients w_j^* set to 0 in the optimal solution of dual-LLP.) □

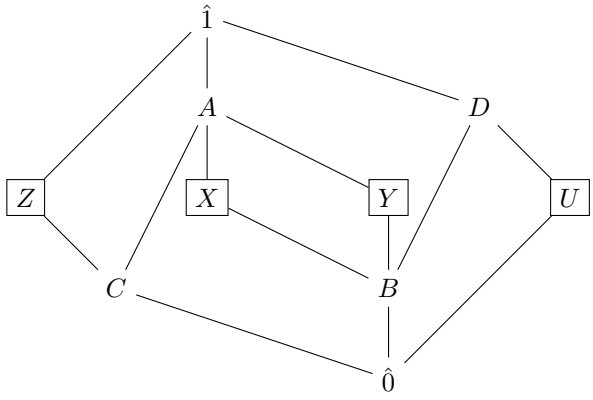

Figure 7: A lattice with an SM-proof that is not good

Example 5.29. Not every SM-proof sequence is good. Consider the following SM-proof sequence of $h(X) + h(Y) + h(Z) + h(U) \geqslant 2h(\hat{1}) + 2h(\hat{0})$ on the lattice in Fig. 7:

$$\begin{aligned} h(X) + h(Y) &\geqslant h(A) + h(B) \\ h(A) + h(Z) &\geqslant h(\hat{1}) + h(C) \\ h(B) + h(U) &\geqslant h(D) + h(\hat{0}) \\ h(C) + h(D) &\geqslant h(\hat{1}) + h(\hat{0}) \end{aligned}$$

From the first three SM-steps, $\mathsf{Labels}(B) = \{2\}$ and $\mathsf{Labels}(C) = \{3\}$, and $\mathsf{Labels}(D) = \{2\}$. Thus, at the last step $\mathcal{A}(C, D) = \varnothing$ which is not good. The inequality admits a different proof sequence, which is good: $(X, Z) \to (C, \hat{1})$, $(Y, U) \to (\hat{0}, D)$, $(C, D) \to (\hat{0}, \hat{1})$. It is unknown whether every SM-proof sequence can be transformed into a good one.

Example 5.30. For the query in Fig. 8, an SM-proof is

$$\begin{aligned} h(X) + h(Y) &\geqslant h(C) + h(A) \\ h(Z) + h(W) &\geqslant h(D) + h(B) \\ h(A) + h(D) &\geqslant h(\hat{1}) + h(\hat{0}) \\ h(B) + h(C) &\geqslant h(\hat{1}) + h(\hat{0}). \end{aligned}$$

After the second SM-step, $\mathsf{Labels}(C) = \{1, 3\}$, $\mathsf{Labels}(D) = \{1, 2\}$, $\mathsf{Labels}(A) = \{2, 3\}$, $\mathsf{Labels}(B) = \{3\}$. Hence, labels 2 and 3 are pushed to copies of $\mathsf{Labels}(\hat{1})$, but label 1 is not present in any set $\mathsf{Labels}(\hat{1})$. So this proof sequence is no good, for a different reason from the previous example.

Here, we show that some lattices don't have any SM-proofs:

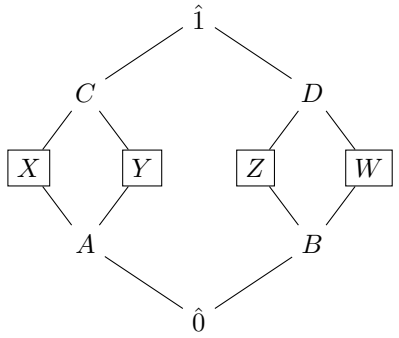

Figure 8: Another example with a bad SM-proof

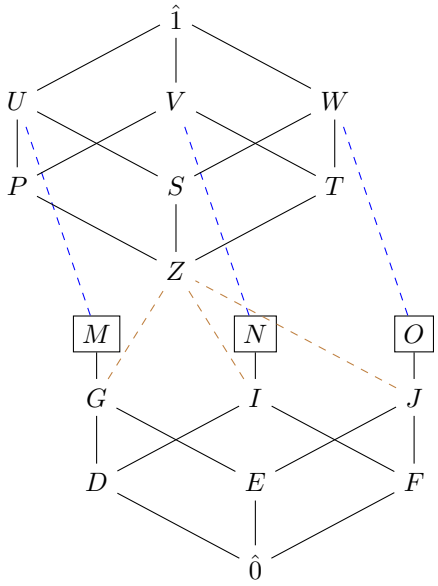

Figure 9: A lattice with no SM-proof sequence

Example 5.31. Not every lattice admits an SM-proof. Consider the lattice in Fig. 9. It satisfies the inequality $h(M) + h(N) + h(O) \geqslant 2h(\hat{1}) + h(\hat{0})$, but there is no SM-proof sequence that derives it. This is somewhat surprising, because Lemma 3.9 showed that every inequality can be proven by adding up submodularity inequalities, and indeed our inequality can be obtained by summing up:

$$h(M) + h(Z) \geqslant h(U) + h(G) \tag{19}$$
$$h(N) + h(Z) \geqslant h(V) + h(I) \tag{20}$$
$$h(O) + h(Z) \geqslant h(W) + h(J) \tag{21}$$
$$h(U) + h(V) \geqslant h(\hat{1}) + h(P) \tag{22}$$
$$h(W) + h(P) \geqslant h(\hat{1}) + h(Z) \tag{23}$$
$$h(G) + h(I) \geqslant h(Z) + h(D) \tag{24}$$
$$h(J) + h(D) \geqslant h(Z) + h(\hat{0}) \tag{25}$$

(The last two inequalities each consists of 2 SM-steps, identical to the proof of $h(xy) + h(xz) + h(yz) \geqslant$

$2h(xyz) + h(\hat{0})$ for the triangle query, see Example 3.10.) But this is not an SM-proof, because in the first step we need $h(Z)$, which is only produced later. More surprisingly, the lattice is normal.

5.3 Meeting LLP-Bound with Conditional SM Algorithm – CSMA

Recall that in the chain proof we decompose an input $h(R_j)$ into a telescoping sum $h(R_j) = \sum_{i \in e_j} (h(R_j \wedge C_i) - h(R_j \wedge C_{i-1}))$. If h is entropic, then $h(Y) - h(X) = h(Y|X)$ is the conditional entropy. Hence, the chain proof decomposes an input $h(R_i)$ into a sum of conditional terms. Taking cues from this conditional decomposition, and the SM-steps of the SM-proof, we devise a new type of proof sequence that resolves the dilemma of the query of Example 5.31, depicted in Fig. 9. In what follows, when $X \preceq Y$ we write $h(Y|X)$ as a short hand for $h(Y) - h(X)$. We assume h is a polymatroid.

Example (5.31 Continued). The key issue is that, starting from $h(M) + h(N) + h(O)$ there is no SM-step that will help prove the desired inequality $h(M) + h(N) + h(O) \geqslant 2h(\hat{1}) + h(\hat{0})$. For example, applying $h(M) + h(N) \geqslant h(\hat{1}) + h(D)$ and we are left to show that $h(D) + h(O) \geqslant h(\hat{1}) + h(\hat{0})$, but $D \vee O = W \prec \hat{1}$ and so this is simply impossible. The trick is to apply an SM-step between G and I, which are *below* the input relations, as suggested by (24). To obtain G and I, we *decompose* $h(M)$ into $h(M|G) + h(G)$ and $h(N)$ into $h(N|I) + h(I)$; these decomposition steps get the proof off the ground.

Formally, our proof strategy involves the following *three* basic rules, instead of just the SM-rule as before:

- CD-rule: for $X \prec Y$, $h(Y) \to h(Y|X) + h(X)$.
- CC-rule: for $X \prec Y$, $h(Y|X) + h(X) \to h(Y)$.
- SM-rule: for $A \not\leqslant B$, $h(A) + h(B|A \wedge B) \to h(A \vee B)$.

(CD stands for *conditional decomposition*, CC *conditional composition*, and SM *sub-modularity*.) Together, they are called the *CSM rules* (for *conditional submodularity*). The vision behind CSMA is to show that

(A) every output inequality (7) can be proved using a series of CD, CC, and SM rules, and

(B) each rule can be interpreted combinatorially to become an algorithmic step, and together they constitute an algorithm (called CSMA) that runs in time $\tilde{O}(2^{h^*(\hat{1})})$

Interestingly, we fell short of objective (A) yet were able to achieve a *stronger* result than objective (B)! In particular, with respect to (A) we conjecture that the three rules above form a complete proof system. We only managed to prove a weaker version of the conjecture (Theorem 5.34) which is sufficient for CSMA to work, at the price of a poly-log factor. On the other hand, with respect to (B) we will work with output inequalities that are more general than (7), and with constraints more general than the fd-constraints.

5.3.1 Conditional LLP (CLLP)

A key technical tool that helps realize the vision is a conditional version of LLP, called the *conditional* LLP (CLLP). A nice by-product of CLLP is that we will be able to compute join queries with given degree bounds, which is even more general than join queries with functional dependencies!

Let $\mathcal{P}$ be some set of pairs $(X, Y) \in L^2$ such that $X \prec Y$. Assume for every pair $(X, Y) \in \mathcal{P}$, there is a given non-negative number $n_{Y|X}$ called a log-*degree bound*. The CLLP is defined as:

$$\begin{array}{lcll} \max & h(\hat{1}) & & \\ \text{s.t.} & h(Y) - h(X) \leqslant n_{Y|X} & & \forall (X,Y) \in \mathcal{P} \\ & h(A \vee B) + h(A \wedge B) - h(A) - h(B) \leqslant 0 & & \forall A \not\leqslant B \in L \\ & h(X) - h(Y) \leqslant 0 & & \forall X \prec Y \in L. \end{array}$$

(By default, $h(\hat{0}) = 0$, and $h(X) \geqslant 0$.) In other words, CLLP requires h to be a polymatroid, subject to the log-degree constraints. The cardinality constraints are a special case of the log-degree constraints $h(Y) - h(\hat{0}) \leqslant n_{Y|\hat{0}} = n_Y$. (Intuitively, cardinality constraints are degree bounds of the empty tuple; fd-constraints $X \to Y$ imposes the Y-degree bound of 0 for every X-tuple; and so the degree bounds $h(Y|X) \leqslant n_{Y|X}$ strictly generalizes both cardinality constraints and fd constraints.) The following is obvious:

Proposition 5.32. LLP *is exactly the same as* CLLP *for the special case when* $\mathcal{P} = \big\{(\hat{0}, R_j) \mid R_j \in \mathbf{R}\big\}$.

Furthermore, we can easily use the log-degree constraints to encode input relations with known maximum degree bounds.

We will also need the dual CLLP. Let $c_{Y|X}, s_{A,B}$ and $m_{X,Y}$ denote the (non-negative) dual variables corresponding to the log-degree, sub-modularity, and monotonicity constraints, respectively. For each $Z \in L - \{\hat{0}\}$, define

$$\mathsf{netflow}(Z) \stackrel{\text{def}}{=} \sum_{\substack{X:X\prec Z\\(X,Z)\in\mathcal{P}}} c_{Z|X} - \sum_{\substack{Y:Z\prec Y\\(Z,Y)\in\mathcal{P}}} c_{Y|Z} + \sum_{\substack{A\not\leqslant B\\A\wedge B=Z}} s_{A,B} + \sum_{\substack{A\not\leqslant B\\A\vee B=Z}} s_{A,B} - \sum_{A:A\not\leqslant Z} s_{A,Z} - \sum_{X:X\prec Z} m_{X,Z} + \sum_{Y:Z\prec Y} m_{Z,Y}.$$

Then, the dual CLLP is

$$\begin{array}{lll} \min & \sum_{(X,Y)\in\mathcal{P}} n_{Y|X} c_{Y|X} & \\ \text{s.t.} & \mathsf{netflow}(\hat{1}) \geqslant 1 & \\ & \mathsf{netflow}(Z) \geqslant 0 & \forall Z \in L - \{\hat{1}, \hat{0}\}. \end{array} \tag{26}$$

Example (5.31 Continued). Consider the lattice in Figure 9. Suppose the input relations are $T(M), T(N), T(O)$ (along with FDs imposing the lattice structure),[12] and we are not given any other bounds on degrees/cardinalities, then the CLLP has $\mathcal{P} = \{(\hat{0}, M), (\hat{0}, N), (\hat{0}, O)\}$ where $n_M = n_{M|\hat{0}} = \log_2 |T(M)|$, $n_N = n_{N|\hat{0}} = \log_2 |T(N)|$, and $n_O = n_{O|\hat{0}} = \log_2 |T(O)|$. In this case CLLP is just LLP. If in addition we also knew, for example, an upper bound d on the degree of G in table $T(M)$, then we can extend the CLLP by adding (G, M) to $\mathcal{P}$ where $n_{M|G} = \log_2 d$.

As mentioned earlier, this lattice satisfies the inequality $2h(\hat{1}) \leqslant h(M) + h(N) + h(O)$, which does not admit an SM-proof. In the CLLP, this inequality corresponds to the constraint $h(\hat{1}) \leqslant \frac{n_M + n_N + n_O}{2}$, which results from the following dual solution.

$$\begin{gathered} c_M = c_N = c_O = 1/2, \\ s_{M,Z} = s_{N,Z} = s_{O,Z} = 1/2, \\ s_{U,V} = s_{P,W} = s_{G,I} = s_{D,J} = 1/2, \end{gathered} \tag{27}$$

(where all dual variables that are not specified above are zeros). In this dual solution, $\mathsf{netflow}(\hat{1}) = 1$ while all other lattice nodes have a netflow value of 0. Hence, this solution is feasible.

5.3.2 CSM proof sequence

The analog of output inequality (7) in the conditional world is

$$\sum_{(X,Y)\in\mathcal{P}} c_{Y|X} h(Y|X) \geqslant h(\hat{1}) \tag{28}$$

Identical to Lemma 3.9, we can show that (28) holds for all polymatroids if there are vectors s and m such that (c, s, m) is feasible to the dual-CLLP.

[12] We abuse notation here, when the same notation symbol is used for different input relations. This is to avoid notation cluttering later in the description of the algorithm.

To answer question **(A)**, we prove a "reachability" lemma that helps us construct a CSM proof sequence. Let (c, s, m) be any feasible solution to the dual-CLLP (26). Let $K \subseteq L$ be a set of lattice elements that contains $\hat{0}$. The *conditional closure* of K (with respect to (c, s, m)) is computed from K by repeatedly applying the following two steps: (1) CD-step ("Conditional Decomposition"): if $Y \in K$ and $X \prec Y$ then add X to K, (2) CC-step ("Conditional Composition"): if $X \in K$ and $c_{Y|X} > 0$ then add Y to K.

Lemma 5.33. *For any dual-feasible solution (c, s, m), let $K \subseteq L$ be a set that contains $\hat{0}$, and $\bar{K}$ be its conditional closure. If $\hat{1} \notin \bar{K}$, then there are two lattice elements $A, B \in \bar{K}$ such that $A \vee B \notin \bar{K}$ and $s_{A,B} > 0$.*

Proof. If $\hat{1} \notin \bar{K}$ then $S \stackrel{\text{def}}{=} \sum_{Z \notin \bar{K}} \mathsf{netflow}(Z) \geqslant 1$. If there is no such pair (A, B), then every dual variable $c_{Y|X}, m_{X,Y}, s_{A,B}$ contributes a non-positive amount to the sum S, which is a contradiction. □

Example (5.31 Continued). Consider the dual solution in (27). Let's take $K = \{\hat{0}\}$ as an example. Because $c_M, c_N, c_O > 0$ (where c_M is just an alias for $c_{M|\hat{0}}$ and so on), the conditional closure $\bar{K}$ contains M, N, O due to the CC-steps. After applying CD-steps, $\bar{K}$ includes all lattice elements below M, N, O as well:

$$\bar{K} = \{\hat{0}, D, E, F, G, I, J, M, N, O\}.$$

(At this point $\bar{K}$ is closed.) Applying Lemma 5.33, we can find $G, I \in \bar{K}$ with $s_{G,I} > 0$. If we add $Z = G \vee I$ to K (and $\bar{K}$), we can apply Lemma 5.33 again and find $M, Z \in \bar{K}$ with $s_{M,Z} > 0$. After adding $U = M \vee Z$ to K, we can find $N, Z \in \bar{K}$ with $s_{N,Z>0}$. After adding $V = N \vee Z$, we will find $U, V \in \bar{K}$ and add $\hat{1} = U \vee V$.

The above lemma allows us to constructively prove a weaker inequality than (28), which is our answer to question **(A)** above.

Theorem 5.34. *Let (c, s, m) be an arbitrary feasible solution to the dual-CLLP, where $c_{Y|X} = q_{Y|X}/d$ and $s_{A,B} = q'_{A,B}/d$ are rational numbers. Let $\mathcal{B}$ be a multiset of variables $h(Y|X)$, $(X, Y) \in \mathcal{P}$, where each variable $h(Y|X)$ occurs $4^{|L|} \cdot q_{Y|X}$ times. Then, there is a sequence of CD-, CC-, and SM-rules that transforms $\mathcal{B}$ into another multiset $\bar{\mathcal{B}}$ which contains the variable $h(\hat{1})$. Moreover, in this sequence all occurrences of identical rules are consecutive in the sequence (hence, they can be combined into a single "rule with a multiplicity").*

Proof. We prove the theorem with an algorithm. We start from the set $K = \{\hat{0}\}$, and keep adding elements to it using conditional closure (CC- and CD-) steps until K is conditionally closed; then we add a new element $A \vee B$ with an SM-step using the pair (A, B) found by Lemma 5.33. This process is repeated until $\hat{1} \in K$. With regard to $\mathcal{B}$, initially we will pretend that for each variable $h(Y|X) \in \mathcal{B}$ we have only $q_{Y|X}$ copies of it instead of $4^{|L|} q_{Y|X}$ copies. While adding elements to K, we will also add variables to $\mathcal{B}$ so eventually $\mathcal{B}$ contains $\leqslant 4^{|L|} q_{Y|X}$ of each variable.

We maintain the following invariants throughout the execution of the algorithm:

- For every $X \in K$, there is at least one copy of $h(X)$ in $\mathcal{B}$. To maintain this invariant, whenever we add a new element X to K while $\mathcal{B}$ does not contain any copy of $h(X)$, we apply some conditional closure rules on $\mathcal{B}$ to produce $h(X)$.
- For every copy of $h(Y|X)$ currently in $\mathcal{B}$, there will always remain at least one copy of $h(Y|X)$ after each step of the algorithm (i.e. in all subsequent multisets $\mathcal{B}$). To maintain this invariant, before applying any rule (where we will be losing one copy of each term on the left-hand side of that rule in order to gain one copy of each term on the right-hand side), we duplicate the multiplicities of all terms in the current $\mathcal{B}$, all previous multisets $\mathcal{B}$, and all rules that have been applied previously. (In effect, we re-run the entire history of rule application once.)

The above invariants are initially satisfied. Now we take the conditional closure of K. Whenever X is added to K due to some $Y \in K$ that satisfies $X \prec Y$ (signalling a CD-step), we first check whether $\mathcal{B}$ contains a copy of $h(X)$. If it does, then no further action is needed. If not, we duplicate multiplicities in the current

and all previous $\mathcal{B}$ and all previous rules, and then we apply a CD-rule $h(Y) \to h(X) + h(Y|X)$. Whenever Y is added to K due to $c_{Y|X} > 0$ for some $X \in K$, we first check whether $h(Y)$ is in $\mathcal{B}$ already. If not, we duplicate all multiplicities as before, and then apply a CC-rule $h(X) + h(Y|X) \to h(Y)$.

Now, suppose K is already conditionally closed, we add $A \vee B$ using an SM-step guaranteed by Lemma 5.33. We check whether $h(A \vee B)$ is in $\mathcal{B}$. If it is not, we check whether $h(B|A \wedge B)$ is in $\mathcal{B}$. If it is not, we duplicate all multiplicities and apply a CD-rule $h(B) \to h(A \wedge B) + h(B|A \wedge B)$. Now, duplicate all multiplicities again, and apply an SM-rule $h(A) + h(B|A \wedge B) \to h(A \vee B)$.

For each element that is added to K, we have to duplicate its multiplicity at most twice, and there are at most $|L|$ such elements.

Finally, we show that the same rule cannot be applied multiple times. In the above, before we applied any CC-rule that produced $h(Y)$, we checked whether $h(Y)$ was already in $\mathcal{B}$. Only if it was not, we applied the rule adding $h(Y)$ to $\mathcal{B}$, and letting the second invariant preserve it in $\mathcal{B}$. The same holds for SM-rules. In the above arguments, we applied CD-rules of the form $h(Y) \to h(X) + h(Y|X)$ in two different places: In the first, we checked that $h(X)$ was not in $\mathcal{B}$ before we added both $h(X)$ and $h(Y|X)$ to $\mathcal{B}$. In the second, we checked that $h(B|A \wedge B)$ was not in $\mathcal{B}$ before we added both $h(B|A \wedge B)$ and $h(A \wedge B)$ to $\mathcal{B}$. □

The series of CC-, CD-, SM-rules with multiplicities is called a *CSM proof sequence*. Interpreted integrally, we think of the proof sequence as having $D \leqslant 4^{|L|}d$ copies of $h(\hat{1})$ that it tries to reach, but at $\bar{\mathcal{B}}$ it reaches at least one copy and we stop.

Example (5.31 Continued). Now, we simulate the proof of Theorem 5.34 on the dual solution in (27) (which corresponds to the inequality $h(\hat{1}) \leqslant \frac{h(M)+h(N)+h(O)}{2}$). We will start off with four copies of the right-hand side (i.e. $2h(M) + 2h(N) + 2h(O)$), and generate a CSM proof sequence that will produce one copy of $h(\hat{1})$ (out of 4 copies).

$$2h(M) \to 2h(M|G) + 2h(G) \quad \text{(adding } G \text{ to } \bar{K} \text{ based on } M \in \bar{K}) \tag{29}$$
$$2h(N) \to 2h(N|I) + 2h(I) \quad \text{(adding } I \text{ to } \bar{K} \text{ based on } N \in \bar{K}) \tag{30}$$
$$2h(I) \to 2h(I|D) + 2h(D) \quad \text{(extracting } h(I|D) \text{ for the next SM-rule)} \tag{31}$$
$$2h(G) + 2h(I|D) \to 2h(Z) \quad \text{(SM-rule based on } s_{G,I} > 0) \tag{32}$$
$$h(Z) + h(M|G) \to h(U) \quad \text{(SM-rule based on } s_{M,Z} > 0) \tag{33}$$
$$h(Z) + h(N|I) \to h(V) \quad \text{(SM-rule based on } s_{N,Z} > 0) \tag{34}$$
$$h(V) \to h(V|P) + h(P) \quad \text{(extracting } h(V|P) \text{ for the next SM-rule)} \tag{35}$$
$$h(U) + h(V|P) \to h(\hat{1}) \quad \text{(SM-rule based on } s_{U,V} > 0) \tag{36}$$

Note that rules (29)…(32) above had multiplicities of 2, because we needed to produce two copies of $h(Z)$: one for (33) and another for (34).

5.3.3 The CSM algorithm (CSMA)

CSMA is our answer to question **(B)**. The algorithm takes as input a join query with functional dependencies *and* maximum degree bounds (if any) from input relations. This input is represented by the set $\mathcal{P}^{(0)}$, corresponding log-degree bounds, and the linear program $\mathsf{CLLP}^{(0)}$. For example, if we use CSMA for the original join query with functional dependencies (with no other max-degree bounds), then we would be starting with $\mathcal{P}^{(0)} = \{(\hat{0}, R_j) \mid R_j \in \mathbf{R}\}$; in this case $\mathsf{CLLP}^{(0)}$ is equivalent to LLP.

Let $h^{(0)}$ and $(c^{(0)}, s^{(0)}, m^{(0)})$ be a pair of primal and dual optimal solutions to $\mathsf{CLLP}^{(0)}$, and OPT be its optimal objective value. CSMA takes the CSM-proof sequence for $\mathsf{CLLP}^{(0)}$ as symbolic *instructions*. For each instruction, CSMA does some computation, spawning a number of sub-problems, creating new intermediate tables for the sub-problems if needed. The final output is contained in the union of outputs of the sub-problems.

For each sub-problem, CSMA constructs a new pair-set $\mathcal{P}'$, and a new linear program CLLP' with dual-feasible solution (c', s', m'). Note that to construct CLLP', there has to be a corresponding log-degree bound $n'_{Y|X}$ for each pair $(X, Y) \in \mathcal{P}'$. The following two invariants are maintained:

(Inv1) For any $(X, Y) \in \mathcal{P}'$, there is a table $T(Y)$ (an input relation of the sub-problem) that "guards" the constraint $h(Y|X) \leqslant n'_{Y|X}$ of CLLP' in the sense that

$$\max_{v \in \Pi_X(T)} \log_2 \deg_T(v) \leqslant n'_{Y|X}.$$

(Note that if $X = \hat{0}$ then the above says $\log_2 |T| \leqslant n'_{Y|\hat{0}} = n'_Y$, which is a cardinality constraint.)

(Inv2) (c', s', m') is feasible to dual-CLLP' with objective value satisfying $\mathsf{OBJ}' \stackrel{\text{def}}{=} \sum_{(X,Y)\in\mathcal{P}} n'_{Y|X} c'_{Y|X} \leqslant \mathsf{OPT}$.

It can be verified that the two invariants are satisfied at $\mathsf{CLLP}^{(0)}$. Next we describe how CSMA deals with each instruction from the CSM sequence.

(1) **CD-rule** $h(Y) \to h(Y|X) + h(X)$ with multiplicity t.

Recall that $h(Y)$ is just a short hand for $h(Y|\hat{0})$. By (Inv1) there is a table $T(Y)$ with $\log_2 |T| \leqslant n_Y$. Define

$$\begin{aligned} n'_{Y|X} &\stackrel{\text{def}}{=} \max_{v \in \Pi_X(T)} \log_2 \deg_T(v) \\ n'_X &\stackrel{\text{def}}{=} \log_2 |\Pi_X(T)| \end{aligned}$$

Lemma 5.35 shows that T can be partitioned into at most $\ell = 2 \log N$ sub-tables $T^{(1)}, \ldots, T^{(\ell)}$ such that $n^{(j)}_{Y|X} + n^{(j)}_X \leqslant n_Y$, for all $j \in [\ell]$, where

$$\begin{aligned} n^{(j)}_{Y|X} &\stackrel{\text{def}}{=} \max_{v \in \Pi_X(T^{(j)})} \log_2 \deg_{T^{(j)}}(v), \\ n^{(j)}_X &\stackrel{\text{def}}{=} \log_2 |\Pi_X(T^{(j)})|. \end{aligned}$$

For each of these sub-tables $T^{(j)}$ of T, we create a sub-problem with T replaced by $T^{(j)}$. For the jth sub-problem, we add $(\hat{0}, X)$ and (X, Y) to $\mathcal{P}$ with corresponding log-degree constraints $h(X) \leqslant n^{(j)}_X$ and $h(Y|X) \leqslant n^{(j)}_{Y|X}$, respectively. We compute the projection of $T^{(j)}$ onto X so we have guards for the two new constraints. Set

$$\begin{aligned} c'_{Y|X} &= c_{Y|X} + t/D, \\ c'_X &= c_X + t/D, \\ c'_Y &= c_Y - t/D \geqslant 0. \end{aligned}$$

If $c'_Y = 0$, then we remove $(\hat{0}, Y)$ from $\mathcal{P}$. By examining $\mathsf{netflow}(Z)$ at each node, we can verify that (c', s, m) is a feasible solution to the new CLLP' with a *reduction* in objective value of $(n_Y - n'_{Y|X} - n'_X)t/D$.

(2) **CC-rule** $h(Y|X) + h(X) \to h(Y)$ with multiplicity t.

Let R be the guard for $h(Y|X) \leqslant n_{Y|X}$ and S for $h(X) \leqslant n_X$. Let θ be a threshold to be determined. If $n_{Y|X} + n_X \leqslant \mathsf{OPT} + \theta$, then we can compute the table $T(Y) \stackrel{\text{def}}{=} S(X) \bowtie R(Y)$ by going over all tuples in S and expanding them using matching tuples in R. The runtime is $\tilde{O}(2^{n_X + n_{Y|X}}) = \tilde{O}(2^{\mathsf{OPT}+\theta})$. The dual solution is modified by setting

$$\begin{aligned} c'_{Y|X} &= c_{Y|X} - t/D, \\ c'_X &= c_X - t/D, \\ c'_Y &= c_Y + t/D, \\ n'_Y &= n_{Y|X} + n_X. \end{aligned}$$

If $n_{Y|X} + n_X > \mathsf{OPT} + \theta$, we will start afresh from an optimal solution to the *current* CLLP. Lemma 5.36 below shows that the current CLLP has an optimal objective value at most $\mathsf{OPT} - \theta/(D-1)$.

(3) **SM-rule** $h(A) + h(B|A \wedge B) \to h(A \vee B)$ with multiplicity t.

This is similar to case (2), and becomes identical to case (2) when $A = A \wedge B = X$ and $B = Y$. Let R be the guard for $h(A) \leqslant n_A$ and S for $h(B|A \wedge B) \leqslant n_{B|A\wedge B}$. If $n_A + n_{B|A\wedge B} \leqslant \mathsf{OPT} + \theta$, then we can compute $T(A \vee B) = R \bowtie S$ in time $\tilde{O}(2^{\mathsf{OPT}+\epsilon})$. If $n_A + n_{A|A\wedge B} > \mathsf{OPT} + \theta$, then we start afresh from a new optimal solution to the current CLLP. The minor difference to case (2) is that we have to modify the variable $s'_{A,B} = s_{A,B} - t/D$. By selecting the correct threshold θ, we can prove that CSMA runs in time $O((\log N)^e 2^{\mathsf{OPT}})$, where e is a data-independent constant (Theorem 5.37).

Example (5.31 Continued). Given a conjunctive query Q whose functional dependencies correspond to the lattice in Figure 9 and whose input relations are $T(M), T(N), T(O)$, where $|T(M)| = |T(N)| = |T(O)| = N$. In the CLLP, we have $\mathcal{P} = \{(\hat{0}, M), (\hat{0}, N), (\hat{0}, O)\}$ where $n_M = n_N = n_O = \log_2 N =: n$. The optimal objective value is $\mathsf{OPT} = \frac{3n}{2}$ (which implies that $|Q| \leqslant N^{\frac{3}{2}}$), and the feasible dual solution given by (27) is optimal. Consider the CSM sequence (29)…(36) that was constructed earlier for (27). We will explain how to run CSMA on this sequence in order to answer Q in time within a polylogarithmic factor of $2^{\mathsf{OPT}} = N^{\frac{3}{2}}$.

- The first rule (29) in the sequence is a CD-rule: $h(M) \to h(M|G) + h(G)$. The corresponding algorithmic step would be to project $T(M)$ on G while making sure that the projection size times the maximum degree of the projection in $T(M)$ does not exceed $|T(M)|$ (i.e. while making sure that the sum of $n'_G \stackrel{\text{def}}{=} \log_2 |\Pi_G(T(M))|$ and $n'_{M|G} \stackrel{\text{def}}{=} \max_{v \in \Pi_G(T(M))} \log_2 \deg_{T(M)}(v)$ does not exceed n_M). If all tuples v in the projection $\Pi_G(T(M))$ have the same degree $\deg_{T(M)}(v)$ (i.e., if $T(M)$ is "uniform" with respect to G), then the required condition is met. Otherwise, let's assume for simplicity that all degrees $\deg_{T(M)}(v)$ are powers of 2 (and $|T(M)|$ is also a power of 2). If this is the case, then based on $\deg_{T(M)}(v)$ we can partition $T(M)$ into a logarithmic number (namely $n_M + 1$) of parts $T^{(0)}(M), T^{(1)}(M), \ldots$ such that the required condition is met in each one of them. (In particular, let $T^{(i)}(M)$ satisfy $n^{(i)}_{M|G} \leqslant i$ and $n^{(i)}_G \leqslant n_M - i$.) Now, the execution of CSMA will split into a logarithmic number of branches, each of which will continue on a different part of $T(M)$. For some arbitrarily-fixed i, let's track the execution of the i-th branch (i.e. the one on $T^{(i)}(M)$).
- The second rule (30) is another CD-rule: $h(N) \to h(N|I) + h(I)$. Similar to above, it will result in the partitioning of $T(N)$ into $O(\log N)$ parts $T^{(0)}(N), T^{(1)}(N), \ldots$ (such that $T^{(j)}(N)$ satisfies $n^{(j)}_{N|I} \leqslant j$ and $n^{(j)}_I \leqslant n_N - j$.) The current i-th branch on $T^{(i)}(M)$ will now branch further into $O(\log N)$ branches corresponding to $T^{(0)}(N), T^{(1)}(N), \ldots$ Let's keep track of the j-th branch (on $T^{(j)}(N)$) for some arbitrary j.
- The third rule (31) is yet another CD-rule: $h(I) \to h(I|D) + h(D)$. It will result in the partitioning of $T(I) \stackrel{\text{def}}{=} \Pi_I(T^{(j)}(N))$ (where $\Pi_I(T^{(j)}(N))$ resulted from the previous CD-step) into $O(\log N)$ parts: $T^{(0)}(I), T^{(1)}(I), \ldots$ (such that $T^{(k)}(I)$ satisfies $n^{(k)}_{I|D} \leqslant k$ and $n^{(k)}_D \leqslant n_I - k$.) Let's track the k-th branch (on $T^{(k)}(I)$).
- The fourth rule (32) is an SM-rule: $h(G) + h(I|D) \to h(Z)$. The corresponding algorithmic step is to join the table $T(G)$ (or more precisely $\Pi_G(T^{(i)}(M))$ from Step 1) with the table $T(I)$ (or more precisely $T(I)$ from Step 3) in order to get a relation $T(Z)$. The time required to compute this join is bounded by the size of $T(G)$ times the maximum degree of D in $T(I)$. Depending on the current execution branch (i, j, k), this time might or might not exceed our budget of $2^{\mathsf{OPT}} = N^{3/2}$. For example, in the branch where both i and j are maximal, we will have $|T(G)| \leqslant 1$ and $|T(I)| \leqslant 1$, hence the join takes $O(1)$ time. On the other hand, in the branch where $i = j = 0$ and k is maximal, each one of $|T(G)|$, $|T(I)|$, and the maximum degree of D in $T(I)$ could be as large as N, hence the join could take time $O(N^2)$. Luckily, Lemma 5.36 below implies that "when one door closes, another one opens": In

those particular branches where the join cannot be computed within our time budget, there are extra constraints that if considered in the CLLP they would reduce the optimal objective value OPT. For example, in the branch where $i = j = 0$ and k is maximal, the value of D is already fixed. This is because the degree of D in table $T(I)$ is maximal, hence it is equal to $|T(I)|$.

- The remaining rules (33)…(36) are similar to the previous ones.

5.3.4 Analysis of CSMA

Lemma 5.35. *Let $X \prec Y$ be in the lattice L. Let $T(Y)$ be a table with $\log_2 |T| \leqslant n_Y$. Then, T can be partitioned into at most $\ell = 2 \log N$ sub-tables $T^{(1)}, \ldots, T^{(\ell)}$ such that $n^{(j)}_{Y|X} + n^{(j)}_X \leqslant n_Y$, for all $j \in [\ell]$, where*

$$\begin{aligned} n^{(j)}_{Y|X} &\stackrel{\text{def}}{=} \max_{v \in \Pi_X(T^{(j)})} \log_2 \deg_{T^{(j)}}(v) \\ n^{(j)}_X &\stackrel{\text{def}}{=} \log_2 |\Pi_X(T^{(j)})|. \end{aligned}$$

Proof. To obtain the copies $T^{(j)}$, observe that the number of elements $v \in \Pi_X(T)$ with log-degree in the interval $[j, j+1)$ is at most $|T|/2^j \leqslant 2^{n_Y - j}$. Hence, if we partition T based on which of the buckets $[j, j+1)$ the log-degree falls into, we would almost have the required inequality: $n^{(j)}_{Y|X} + n^{(j)}_X \leqslant (j+1) + (n_Y - j) = n_Y + 1$. To resolve the situation, we partition each $T^{(j)}$ into two equal-sized tables. Overall, we need $\ell = 2 \log N$. □

Lemma 5.36. *Given a CLLP whose optimal objective value is OPT, and a feasible dual solution (c, s, m) whose objective value is $\mathsf{OBJ} = \sum_{(\bar{X}, \bar{Y}) \in \mathcal{P}} c_{\bar{Y}|\bar{X}} n_{\bar{Y}|\bar{X}}$, let $c_Y > 0$ for some $(\hat{0}, Y) \in \mathcal{P}$.*

- *If $n_Y > \mathsf{OBJ}$, then $\mathsf{OPT} < \mathsf{OBJ}$.*
- *Given $\theta \geqslant 0$, $0 < \epsilon < 1$, $\epsilon \leqslant c_Y$, if $n_Y > \mathsf{OBJ} + \theta$, then $\mathsf{OPT} < \mathsf{OBJ} - \frac{\epsilon\theta}{1-\epsilon}$.*

Proof. Let (c', s, m') be a dual solution obtained by setting $c'_Y = c_Y - \epsilon$, $m'_{Y,\hat{1}} = m_{Y,\hat{1}} + \epsilon$ and keeping other c' and m' values identical to their c and m counterparts. (c', s, m') is not necessarily feasible. However, it satisfies $\mathsf{netflow}(\hat{1}) \geqslant 1 - \epsilon$ and $\mathsf{netflow}(Z) \geqslant 0$ for all $Z \in L - \{\hat{1}, \hat{0}\}$. Let (c'', s'', m'') be a dual solution obtained by multiplying (c', s, m') by $\frac{1}{1-\epsilon}$. Now (c'', s'', m'') is indeed feasible. Let OBJ'' be its objective value. Because it is dual feasible, $\mathsf{OBJ}'' \geqslant \mathsf{OPT}$.

$$\begin{aligned} \mathsf{OBJ}'' &= \sum_{(\bar{X}, \bar{Y}) \in \mathcal{P}} c''_{\bar{Y}|\bar{X}} n_{\bar{Y}|\bar{X}} \\ &= \sum_{(\bar{X}, \bar{Y}) \in \mathcal{P}} \frac{c'_{\bar{Y}|\bar{X}}}{1-\epsilon} n_{\bar{Y}|\bar{X}} \\ &= \frac{1}{1-\epsilon} \sum_{(\bar{X}, \bar{Y}) \in \mathcal{P}} c_{\bar{Y}|\bar{X}} n_{\bar{Y}|\bar{X}} - \frac{\epsilon}{1-\epsilon} n_Y \\ &< \frac{1}{1-\epsilon} \mathsf{OBJ} - \frac{\epsilon}{1-\epsilon} (\mathsf{OBJ} + \theta) \\ &= \mathsf{OBJ} - \frac{\epsilon\theta}{1-\epsilon}. \end{aligned}$$

□

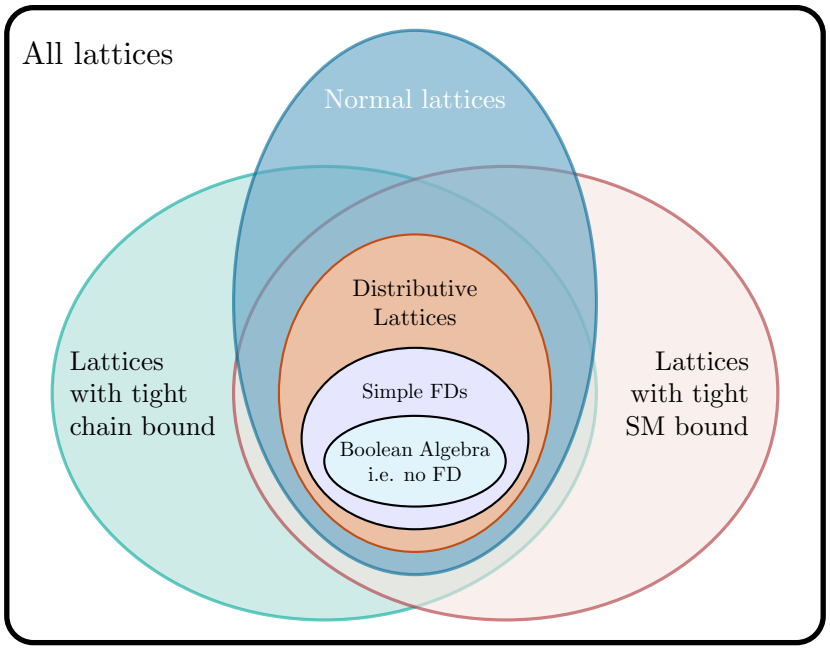

Figure 10: A summary of the lattices discussed in this paper

Theorem 5.37. *CSMA runs in time $O(N + (\log N)^e 2^{\mathsf{OPT}})$, where OPT is the optimal objective value for $\mathsf{CLLP}^{(0)}$, N is the input size, and e is a data-independent constant.*

Proof. Since the number of CSM rules is at most $|L|^2$ (these are rules with multiplicities), so no rule is repeated. The worst case is obtained when the algorithm branches as far as possible only to have to reduce the objective by $\epsilon = \frac{\theta}{D-1}$ at the very end and all leaf nodes of the execution tree have to be started afresh with the new optimal value. Let $x = 2^{\mathsf{OPT}}$ and $T(x)$ denote the runtime of the algorithm on the $\mathsf{CLLP}^{(0)}$ with objective value OPT. Then, the recurrence for the runtime is

$$\begin{aligned} T(x) &= (\ell)^{|L|^2} T\left(2^{\mathsf{OPT}-\frac{\theta}{D-1}}\right) + \ell^{|L|^2} 2^\theta x \\ &= (\ell)^{|L|^2} T\left(\frac{x}{2^{\theta/(D-1)}}\right) + \ell^{|L|^2} 2^\theta x \end{aligned}$$

To get the exponential decay effect we set θ so that $2^{\theta/(D-1)} = 2\ell^{|L|^2}$, which means $\theta = (D-1)(|L|^2 \log_2 \ell + 1)$. □

6 Conclusions

We studied ways to prove worst-case output size bounds, and algorithms meeting the bounds for join queries with functional dependencies. A main aim was to design an algorithm running within time bounded by the entropy-based linear program proposed by Gottlob et al. [11]. For this purpose we developed new, lattice theoretic techniques, of independent interest. We described several classes of lattices, and several ways to prove upper or lower bounds on the worst-case query output, summarized in Fig. 10. On the algorithmic side, we devised the novel idea of turning a proof of an inequality into an algorithm. Three proof techniques lead to three different algorithms with increasing complexity. Our main algorithm does meet GLVV bound, up to a poly-log factor. The algorithm manages to solve a stronger problem, where input relations have prescribed maximum degree bounds, of which functional dependencies and cardinality bounds are special cases.

References

[1] Aberger, C. R., Nötzli, A., Olukotun, K., and Ré, C. Emptyheaded: Boolean algebra based graph processing. *CoRR abs/1503.02368* (2015).

[2] Aref, M., ten Cate, B., Green, T. J., Kimelfeld, B., Olteanu, D., Pasalic, E., Veldhuizen, T. L., and Washburn, G. Design and implementation of the logicblox system. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, Melbourne, Victoria, Australia, May 31 - June 4, 2015* (2015), T. K. Sellis, S. B. Davidson, and Z. G. Ives, Eds., ACM, pp. 1371–1382.

[3] Atserias, A., Grohe, M., and Marx, D. Size bounds and query plans for relational joins. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA* (2008), pp. 739–748.

[4] Balister, P., and Bollobás, B. Projections, entropy and sumsets. *Combinatorica 32*, 2 (2012), 125–141.

[5] Benedikt, M., Leblay, J., and Tsamoura, E. Querying with access patterns and integrity constraints. *PVLDB 8*, 6 (2015), 690–701.

[6] Boyd, S., and Vandenberghe, L. *Convex Optimization*. Cambridge University Press, 2004.

[7] Chu, S., Balazinska, M., and Suciu, D. From theory to practice: Efficient join query evaluation in a parallel database system. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, Melbourne, Victoria, Australia, May 31 - June 4, 2015* (2015), pp. 63–78.

[8] Chung, F. R. K., Graham, R. L., Frankl, P., and Shearer, J. B. Some intersection theorems for ordered sets and graphs. *J. Combin. Theory Ser. A 43*, 1 (1986), 23–37.

[9] Demetrovics, J., Libkin, L., and Muchnik, I. B. Functional dependencies in relational databases: A lattice point of view. *Discrete Applied Mathematics 40*, 2 (1992), 155–185.

[10] Gogacz, T., and Toruńczyk, S. Entropy bounds for conjunctive queries with functional dependencies. *CoRR abs/1512.01808* (2015).

[11] Gottlob, G., Lee, S. T., Valiant, G., and Valiant, P. Size and treewidth bounds for conjunctive queries. *J. ACM 59*, 3 (2012), 16.

[12] Grohe, M. Bounds and algorithms for joins via fractional edge covers. In *In Search of Elegance in the Theory and Practice of Computation*. Springer Berlin Heidelberg, 2013, pp. 321–338.

[13] Grohe, M., and Marx, D. Constraint solving via fractional edge covers. In *SODA* (2006), ACM Press, pp. 289–298.

[14] Harremoës, P. Functional dependences and Bayesian networks. In *Proceedings WITMSE 2011* (Helsinki, 2011).

[15] Joglekar, M., and Ré, C. It's all a matter of degree: Using degree information to optimize multiway joins. *CoRR abs/1508.01239* (2015).

[16] Levene, M. A lattice view of functional dependencies in incomplete relations. *Acta Cybern. 12*, 2 (1995), 181–207.

[17] Marx, D. Tractable hypergraph properties for constraint satisfaction and conjunctive queries. *J. ACM 60*, 6 (2013), Art. 42, 51.

[18] Ngo, H. Q., Porat, E., Ré, C., and Rudra, A. Worst-case optimal join algorithms: [extended abstract]. In *PODS* (2012), pp. 37–48.

[19] Ngo, H. Q., Ré, C., and Rudra, A. Skew strikes back: new developments in the theory of join algorithms. *SIGMOD Record 42*, 4 (2013), 5–16.

[20] Radhakrishnan, J. 6. entropy and counting. *Computational Mathematics, Modelling and Algorithms* (2003), 146.

[21] Schrijver, A. *Combinatorial optimization. Polyhedra and efficiency. Vol. B*, vol. 24 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2003. Matroids, trees, stable sets, Chapters 39–69.

[22] Stanley, R. P. *Enumerative combinatorics. Volume 1*, second ed., vol. 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2012.

[23] Veldhuizen, T. L. Triejoin: A simple, worst-case optimal join algorithm. In *Proc. 17th International Conference on Database Theory (ICDT), Athens, Greece, March 24-28, 2014.* (2014), N. Schweikardt, V. Christophides, and V. Leroy, Eds., OpenProceedings.org, pp. 96–106.

[24] Yeung, R. W. *Information Theory and Network Coding*, 1 ed. Springer Publishing Company, Incorporated, 2008.

[25] Zhang, Z., and Yeung, R. W. On characterization of entropy function via information inequalities. *IEEE Transactions on Information Theory 44*, 4 (1998), 1440–1452.

A Additional Material for Sec 1

We show that the query shown in (2) has output size bounded by $O(N^{3/2}, Nd_1, Nd_2)$. W.L.O.G we may allow an outgoing edge to be colored with more than one color, i.e. the FD $xy \to c_1$ is not required. This is because the largest output is obtained when the number of pairs (x, y) in R is maximized, while $|R| \leqslant N$, and this happens when each edge has only one color. The bound follows easily because

$$h(zx) + h(c_1) \geqslant h((zxc_1)^+) = h(xyzc_1c_2).$$

B Additional Material for Sec 3

The following proposition can be found in [21], pp. 774; we reproduce it here for completeness.

Proposition B.1 (Lovasz's monotonization)**.** *Let $\mathbf{L} = (L, \preceq)$ be a lattice, and h be a non-negative submodular function on the lattice. Define $\bar{h} : L \to \mathbb{R}$ by*

$$\bar{h}(X) = \begin{cases} 0 & X = \hat{0} \\ \min_{Y: X \preceq Y} h(Y) & X \neq \hat{0} \end{cases}$$

Then, $\bar{h}$ is a polymatroid where $\bar{h}(\hat{1}) = h(\hat{1})$ and $\forall X$, $\bar{h}(X) \leqslant h(X)$.

Proof. We verify the only non-trivial property, that $\bar{h}$ is sub-modular. Fix $X \not\preceq Y$, and let $\bar{X} = \operatorname{argmin}_{Z: X \preceq Z} h(Z)$ and $\bar{Y} = \operatorname{argmin}_{Z: Y \preceq Z} h(Z)$. Noting that $X \vee Y \preceq \bar{X} \vee \bar{Y}$ and $X \wedge Y \preceq \bar{X} \wedge \bar{Y}$, we have

$$\begin{aligned} \bar{h}(X) + \bar{h}(Y) &= h(\bar{X}) + h(\bar{Y}) \\ &\geqslant h(\bar{X} \wedge \bar{Y}) + h(\bar{X} \vee \bar{Y}) \\ &\geqslant \bar{h}(X \wedge Y) + \bar{h}(X \vee Y). \end{aligned}$$

□